Project Status Report

High End Computing Capability
Strategic Capabilities Assets Program

March 10, 2018

Dr. Rupak Biswas – Project Manager
NASA Ames Research Center, Moffett Field, CA
Rupak.Biswas@nasa.gov
(650) 604-4411

Engineers Deploy Automounting to Improve Compute Node Reliability

- HECC engineers recently deployed automatic mounting and unmounting of filesystems in the HECC environment. This results in reduced load on file servers, improved system stability, and provides some additional free memory on the compute nodes.
- Filesystems are mounted as requested by either batch jobs or on demand, and are automatically unmounted after a period of inactivity. This scheme reduces the network and server load by only mounting the filesystem as needed.
- The largest benefit of this change is improved system stability, as only the filesystems required are mounted. This means that if a filesystem experiences an issue, only a subset of nodes will be impacted.
- Not mounting all of the filesystems also slightly reduces the amount of memory required on the compute node for caching. The memory saved can instead be utilized by a user's application.

Mission Impact: HECC's implementation of automounting improves the reliability of compute nodes by reducing the impact of filesystem issues, in turn providing a more stable compute environment for the users.

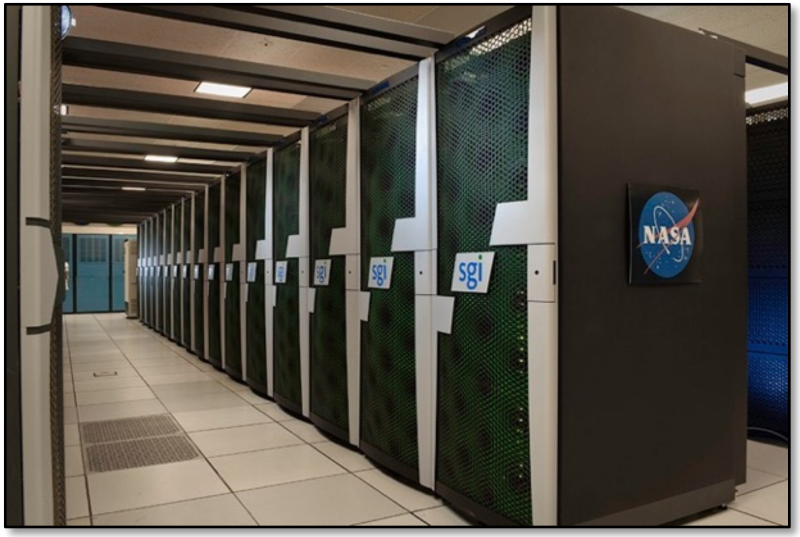

There are 17 Lustre or NFS filesystems that can be mounted on the compute nodes. Automounting reduces the number down to the few that are required to run batch jobs.

POCs: Bob Ciotti, bob.ciotti@nasa.gov, (650) 604-4408, NASA Advanced Supercomputing (NAS) Division;
Davin Chan, davin.chan@nasa.gov, (650) 604-3613, NAS Division, CSRA LLC

Applications Experts Improve CoSMoS Job Efficiency by a Factor of 7

- HECC Applications Performance and Productivity (APP) experts, working with Jessica Lovering of the U.S. Geological Survey, improved her job efficiencies by a factor of 7 by porting the Delft3D suite of modeling codes, developing a set of scripts based on GNU parallel, and tuning OpenMP parallelism.
- The Coastal Storm Modeling System (CoSMoS) comprises two coupled codes: the FLOW module of Delft3D for the hydrodynamics of the water system and the WAVE module, based on the Simulating Waves Nearshore (SWAN) code.
  - Both executables are run simultaneously, and each run takes 10 hours on a single Sandy Bridge node.
  - Both executables are run as single-process jobs, and GNU parallel was used to pack seven model simulations on a single node. Unfortunately, this increased the wall-time to 67 hours.
  - The APP team determined that the huge increase in run time was due to the periodic running of SWAN using 32 threads for each simulation, and with seven simulations running on one node—the 16 cores were overwhelmed by the 32x7 number of OpenMP threads.
  - Limiting the number of threads for each SWAN process to 2 decreased the wall-time to 16 hours. And with 4 threads, the run time dropped to just 9 hours—for a more than 7-fold decrease in SBUs used, compared to running each model simulation on a separate node.

Mission Impact: Computational efficiencies provided by HECC staff enable researchers to study additional simulations of coastal hazards for a wider range of sea-level rise and storm conditions. By successfully packing seven simulations into a single node, HECC improved node utilization and reduced the total node-hours needed—both by a factor of 7.

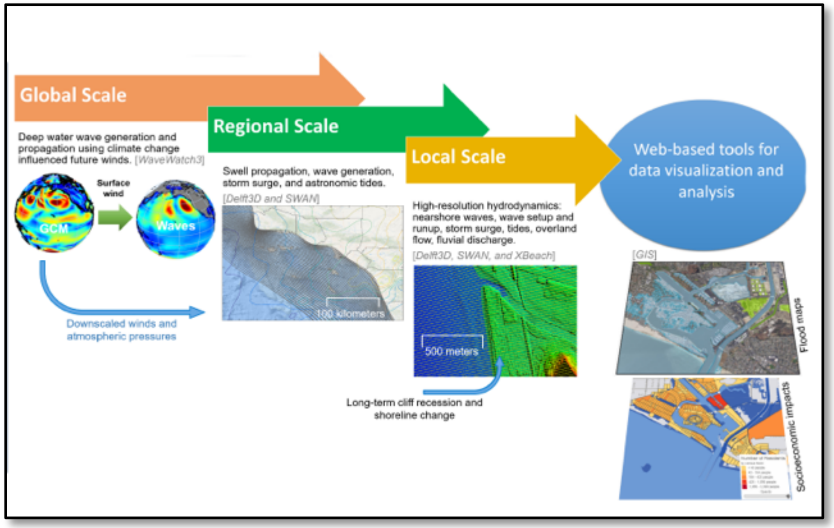

Coastal Storm Modeling System (CoSMoS): Physics-based numerical modeling system for assessing coastal hazards due to climate change.

POC: Johnny Chang, johnny.chang@nasa.gov, (650) 604-4356, NASA Advanced Supercomputing Division, CSRA LLC

Mars Climate Model Data Adapted for Use with hyperwall Visualization Tool

- The HECC visualization team recently started working with scientists at NASA Ames' Mars Climate Modeling Center, and was able to convert the first dataset in less than two weeks, so that it could be used with the multi-movie hyperwall visualization tool (the same tool used with a 1/48th degree ECCO dataset; see the January 2018 report).
- The Mars Climate Modeling group used the HECC tool with their data during two different work sessions, and the visualization team has already implemented two requested features:
  - A zoom and pan capability.
  - Adaption of the tool that produces vertical slices and time-space plots so that it works with the Mars simulation data.
- The "Vis" team expects to soon convert output from additional simulation runs and work with the Mars group in sessions to compare the different runs.

Mission Impact: As a result of swift work by HECC visualization experts, scientists from the Mars Climate Modeling Center are now able to quickly look at a large fraction of their simulation output at once, which will increase the speed of their analyses.

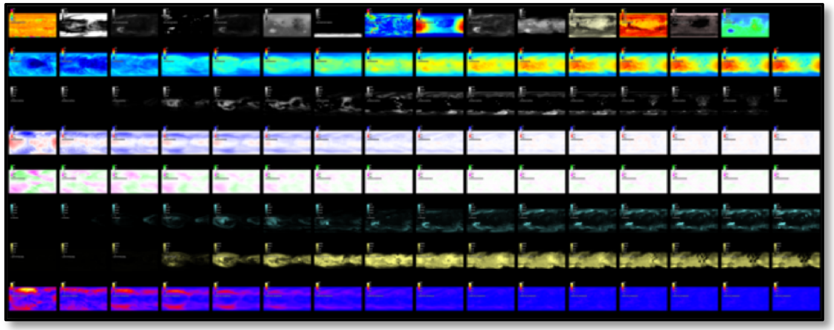

Snapshot of the NASA hyperwall showing Mars climate data. The top row shows the 2D fields, while the next seven rows show the 3D fields for temperature, water ice, zonal wind, meridional wind, dust, specific humidity, and horizontal wind speed. Each row that shows a 3D field displays 16 of the 28 pressure levels, with the highest level on the left and proceeding to the surface on the right. *David Ellsworth, NASA/Ames*

POCs: David Ellsworth, david.ellsworth@nasa.gov, (650) 604-0721, NASA Advanced Supercomputing (NAS) Division, CSRA LLC; Chris Henze, chris.henze@nasa.gov, (650) 604-3959, NAS Division

Tools Team Released myNAS Portal Website to HECC Users

- The HECC Tools team released the myNAS website to all HECC users this month, following the October 2017 release to principal investigators (PIs). This site now provides users with near-real-time information on their running, waiting, held, and recently finished jobs across all their GIDs.
- Job listings can be filtered by users, machines, queues, models, and other parameters. Users can also view utilization bar and donut charts, which show allocation and accounting information at a glance.
- The Tools team collaborated with the Application Performance & Productivity team to develop the myNAS website, which is hosted on the portal.nas.nasa.gov server. Users simply authenticate using Launchpad, with either a NASA Smartcard or RSA token, and all of their job and usage data are retrieved and loaded within a few seconds.
- New features in this release include updates to the finished jobs and usage reports. Finished jobs tables now mark jobs that are not charged to the GID due to queue type or exit status. When a user logs in, they can get a report of their own usage, along with cumulative usage for the rest of the GID members.
- myNAS uses multiple scripts and web services to gather, store, process, and present detailed information on thousands of active and completed jobs, updated for users about every 10 minutes.
- The Tools team is planning several new feature releases of the myNAS portal over the next year.

Mission Impact: The myNAS portal provides users and principal investigators with new tools to analyze and optimize their use of HECC resources.

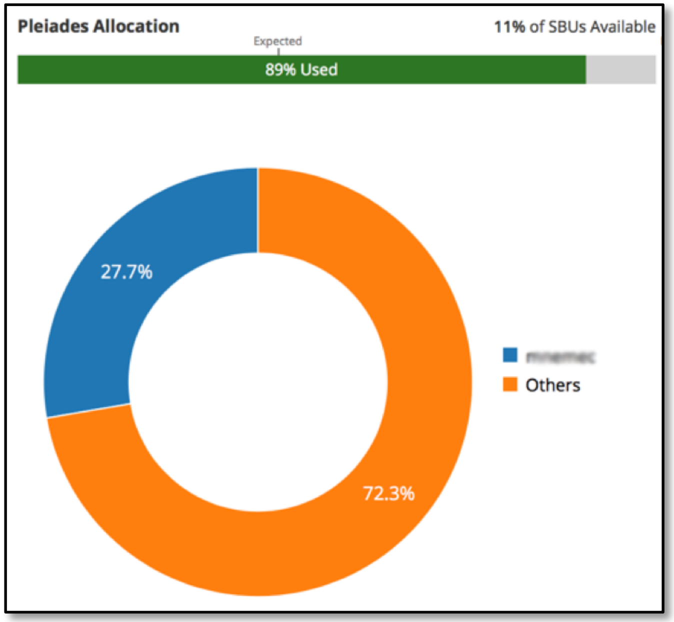

Screenshot from myNAS displaying a bar graph with the percentage of total GID allocation used to date vs. expected (linear) usage. The donut chart indicates the relative amount of SBUs used by the logged-in user against the total for the rest of the GID members. In contrast, PIs see a detailed breakdown of usage by each user.

POCs: John Hardman, john.hardman@nasa.gov, (650) 604-0417, NASA Advanced (NAS) Supercomputing (NAS) Division, CSRA LLC; Ryan Spaulding, ryan.c.spaulding@nasa.gov, (408) 772-6567, NAS Division, ADNET Systems

HECC Develops Collaborative Database of Solar Flares for Heliophysics Research

- As part of HECC work to provide a data portal for collaborative heliophysics research, staff developed the Interactive Multi-Instrument Database of Solar Flares. This multi-instrument database is essential to facilitate modeling and analysis of the impact of solar flare radiation.
- The HECC Big Data team integrated and added security checks to the database in a new heliophysics portal, hosted on the NAS website: http://helioportal.nas.nasa.gov.
  - A web interface allows users to search for uniquely identified flare events based on their physical characteristics and other pre-defined criteria, in order to investigate their radiation properties, including extreme ultraviolet radiation and X-ray radiation.
  - Currently, data from three primary flare lists (National Oceanic and Atmospheric Administration, NASA, and Lockheed Martin) and a variety of other event catalogs from spacecraft and ground-based observations are integrated into the database. The datasets are available starting from 2002, and data is kept current by updating nightly from all sources.

Mission Impact: The heliophysics portal and database enhance research studies in impact of solar flares by providing an integrated database of reported flares and ground-based observations.

The Heliophysics Portal website provides highlights of the latest solar events. This screenshot shows solar flare energy released during a coronal mass ejection. The inset shows the integrated search available across multiple sources, using a wide variety of filters to query the data.

POC: Shubha Ranjan, shubha.ranjan@nasa.gov, (650) 604-1918, NASA Advanced Supercomputing Division

Simulations of Flexible Proteins to Further Origin-of-Life Studies

- Origin-of-life researchers at NASA Ames are running molecular dynamics (MD) protein simulations on Pleiades to help answer questions about the nature of the most primitive organisms.
- From these complex simulations (accompanied by lab experiments run by collaborators) the researchers identified hypothetical ancient protein structures that are shorter and much more flexible than most modern proteins that benefit from billions of years of evolution.
- In their study, the researchers carried out MD simulations of the original protein and its mutants. Surprisingly, two mutations expected to break the structure of the protein and abolish its function did not do so, as confirmed experimentally. This unusual robustness of the protein to mutations arises from small structural rearrangements in the core, facilitated by its flexibility.
- Results from experiments and the simulations led researchers to hypothesize that—contrary to some other scientific assumptions—small, flexible proteins, which were quite different from their modern counterparts, could have carried out functions required for primordial life without extensive evolutionary optimization that conferred rigidity on modern proteins.

HECC provided supercomputing resources and services in support of this work.

Mission Impact: Simulations of flexible proteins, enabled by HECC resources, help develop new hypotheses to support NASA's astrobiology missions, with a goal to understand the origin, evolution, distribution, and future of life in the Universe.

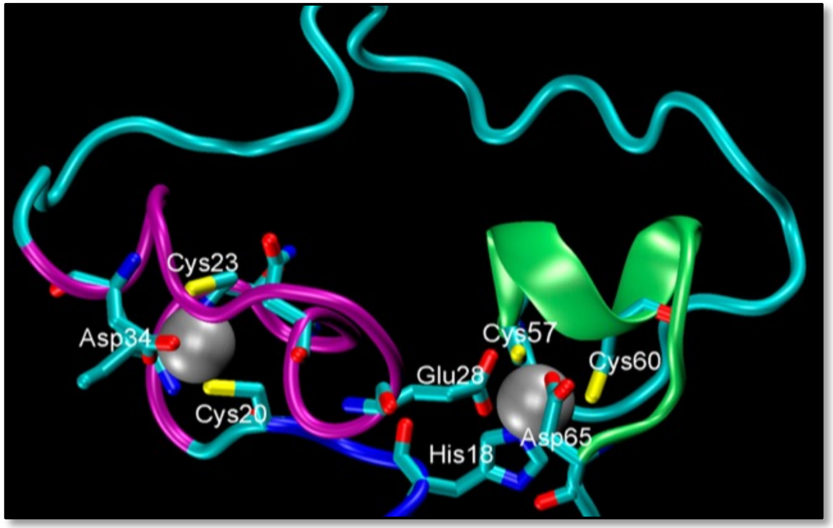

A computer model of an early protein that may be the archetype of primordial enzymes. The protein does not contain conventional elements of the secondary structure characteristic of modern water-soluble proteins, but instead is built of a flexible, catalytic loop supported by a small hydrophilic core containing zinc atoms (grey balls).

POC: Andrew Pohorille, andrew.pohorille@nasa.gov, (650) 604-5759, Exobiology Branch, NASA Ames Research Center

Pleaides Used to Show Impacts of 2015–2016 El Niño on Global Carbon Cycle

- The Science Data Operations (SDOS) team at NASA's Jet Propulsion Laboratory generates data products supporting the Orbiting Carbon Observatory-2 (OCO-2) mission to quantify the exchange of carbon dioxide between Earth's atmosphere, oceans, and plants on land.
- Using existing OCO-2 data, advanced modeling tools, and HECC supercomputers, scientists showed that the 2015-2016 El Niño had a significant impact on the global carbon cycle—Amazon, Africa, and Southwest Asia released ~50% more carbon dioxide than in a typical year.
- To achieve improved precision, it takes about 5 minutes to process each of approximately 200,000 OCO-2 soundings. About 1 million soundings are collected each day, and using the Pleiades supercomputer allows the team to reprocess all cloud-free scenes monthly with updated set calibration coefficients and obtain a larger, higher resolution dataset of about 5 million soundings per month.
- The team also addressed a bias issue in a new version of their code, which was used to reprocess the entire dataset over 10 weeks.

HECC provided supercomputing resources and services in support of this work.

Mission Impact: HECC resources enabled scientists to produce Orbiting Carbon Observatory data products and shortened the time needed to reprocess 34 months of mission data.

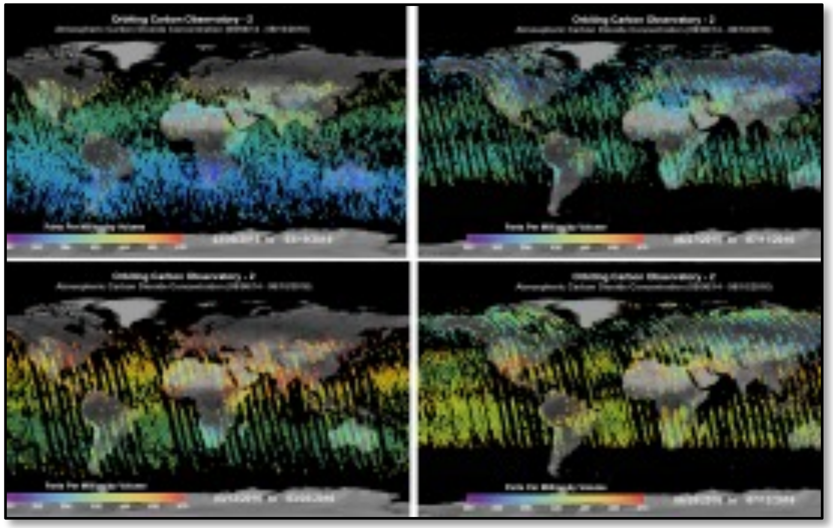

Four maps illustrating the atmospheric carbon dioxide measured from NASA's Orbiting Carbon Observatory-2 (OCO-2) mission. The increase in CO_2 concentration from one year to the next can be seen, as well as the seasonal changes between early spring and summer, where plant growth reduced CO_2 concentrations in July.

POCs: Cecilia Cheng, cecelia.s.cheng@jpl.nasa.gov, (818) 216-6256, AnnMarie Eldering, annmarieeldering@jpl.nasa.gov, (818) 354-4941, Jet Propulsion Laboratory

Simulating NASA's X-57 Electric Concept Aircraft to Predict Aerodynamic Performance

- Aerospace engineers from several NASA centers are running computational fluid dynamics (CFD) simulations of the X-57 electric concept aircraft on Pleiades, using NASA's Launch Ascent and Vehicle Aerodynamics (LAVA) CFD solver.
  - The ultimate objective of this study is to verify that the X-57 design will provide a 5-fold decrease in energy consumption at high-speed cruise.
  - To do this, researchers from NASA's Ames, Armstrong, and Langley Research Centers are using CFD to collaboratively build an aerodynamic database and to predict the aero-propulsive effects of the wing tip propulsors.
  - While experimental wind tunnel data exists for a 19% scale model, CFD tools are needed to adjust this data to the flight Reynolds number and report these modified results. Each research center tested a variety of grid paradigms, mesh fineness levels, and solver algorithms.
- Results show that LAVA can be used to model and simulate the X-57 design's complex geometry, and that the aircraft's aerodynamic performance can be predicted accurately.

HECC provided supercomputing resources and services in support of this work.

Mission Impact: Analyzing the unique X-57 design requires high-quality computational grids and robust numerical algorithms that can handle its complex geometry. These can be studied under NASA's LAVA solver framework using HECC resources.

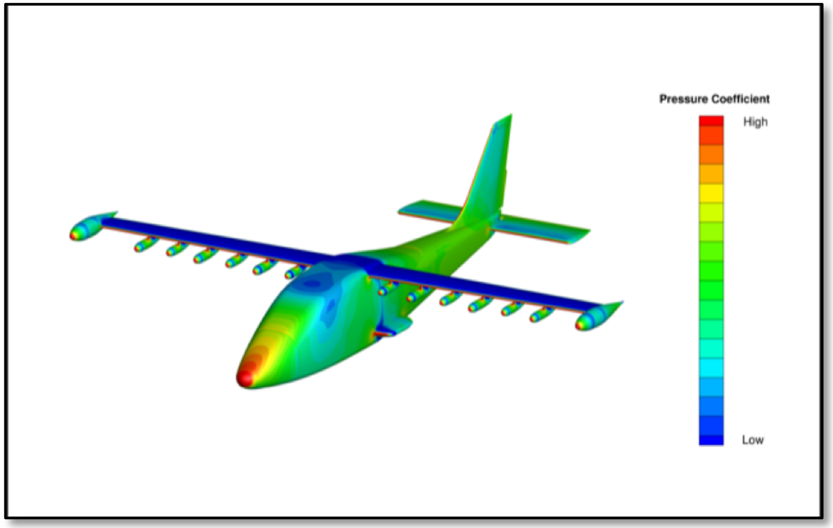

Pressure field over a 19% scale model of the X-57 Modification III design with nominal control surface deflections, at zero-degree angle of attack, inside the 12-ft. subsonic wind tunnel at NASA's Langley Research Center. Quantified simulation results are used to measure solution variation caused by different codes, grids, and solver settings. *Jared Duensing, NASA/Ames*

POCs: Jared Duensing, jared.c.duensing@nasa.gov, (650) 604-4527, NASA Advanced Supercomputing (NAS) Division, Science & Technology Corp.; Cetin Kiris, cetin.c.kiris@nasa.gov, (650) 604-4485, NAS Division

HECC Facility Hosts Several Visitors and Tours in February 2018

- HECC hosted 8 tour groups in February; guests learned about the agency-wide missions being supported by HECC assets, and some groups also viewed the D-Wave 2X quantum computer system. Visitors this month included:
  - Michael Freilich, Earth Science Director in the Science Mission Directorate at NASA Headquarters.
  - Yoav Evenstein, Head of Data Science for Research and Development Innovation for the Israeli Ministry Defense Technical Support.
  - A group of scientists from the United States Geological Survey.
  - Patricia Cornwell, an international best-selling author, who is doing research on quantum computing for an upcoming novel.
  - A group of developers and stake holders who are part of the NASA Human Computers Interaction and Mission Assurance Software face-to-face meeting held here at Ames.
  - 20 junior and senior high school students and instructors from Red Bluff High School, CA who are participating in NASA's Lassen Astrobiology Intern Program.
  - A fourth-grade class from Belmont Oaks Academy, a local school near Ames.

Science managers view images ocean models on the NASA hyperwall at Ames. From left: Rupak Biswas, Director, Ames Exploration Technology Directorate; Jaya Bajpayee, Ames Deputy Director of Science; Michael Freilich, Earth Science Director at NASA Headquarters; Ryan Spackman, Ames Earth Science Division Chief; Piyush Mehrotra, Chief, NAS Division at Ames; and Chris Henze, NAS Visualization Lead.

POC: Gina Morello, gina.f.morello@nasa.gov, (650) 604-4462, NASA Advanced Supercomputing Division

Papers

- "**Connective Dynamics and Disequilibrium Chemistry in the Atmosphere of Giant Planets and Brown Dwarfs**," B. Bordwell, B. Brown, J. Oishi, The Astrophysical Journal, vol. 854, no. 1, February 6, 2018. *
http://iopscience.iop.org/article/10.3847/1538-4357/aaa551/meta
- "**A New Estimate of North American Mountain Snow Accumulation From Regional Climate Models**," M. Wrzesien, et al., Geophysical Research Letters, vol. 45, issue 3, February 10, 2018. *
http://onlinelibrary.wiley.com/doi/10.1002/2017GL076664/full
- "**Quantifying Uncertainty in Discrete-Continuous and Skewed Data with Bayesian Deep Learning**," T. Vandal, E. Kodra, J. Dy. S. Ganguly, R. Nemani, A. Ganguly, arXiv:1802.04742 [cs.LG], February 13, 2018. *
https://arxiv.org/abs/1802.04742
- "**Quantifying the Effect of Non-Larmor Motion of Electrons on the Pressure Tensor**," H. Che, et al., arXiv:1802.05207 [physics.plasm-ph], February 14, 2018. *
https://arxiv.org/abs/1802.05207
- "**The Frequency of Window Damage Caused by Bolide Airbursts: A Quarter Century Case Study**," N. Gi, P. Brown, M. Aftosmis, arXiv:1802.07299 [astro-ph.EP], February 20, 2018. *
https://arxiv.org/abs/1802.07299
- "**Ocean Submesoscales as a Key Component of the Global Heat Budget**," Z. Su, J. Wang, P. Klein, A. Thompson, D. Menemenlis, Nature Communications, vol. 9, February 22, 2018. *
https://www.nature.com/articles/s41467-018-02983-w

** HECC provided supercomputing resources and services in support of this work*

News and Events

- **'Handyman of Proteins' Got Life Started**, *Astrobiology Magazine*, February 15, 2018-– Researchers at NASA Ames Research Center are running supercomputer simulations of hypothetical ancient proteins to better understand the origin and evolution of life in the universe.
  https://www.astrobio.net/alien-life/handyman-proteins-got-life-started/
- **Open Source Powers Supercomputing**, *Communications of the ACM,* February 20, 2018—Open source experts say it comes as little surprise that the top 500 supercomputers in the world are running Linux. The article quotes chief systems architect Bob Ciotti, and gives examples of NASA mission projects run on Pleiades.
- **Student Interns Achieve Dynamic UAS Model Using NASA Supercomputer**, *Bay Area Environmental Research Institute*, February 23, 2018—Students from Cal Poly San Luis Obispo and the University of Washington used the Pleiades supercomputer to create simulations of an unmanned autonomous system aircraft that will support a variety of aerodynamic analyses in support of future science missions.
  https://baeri.org/baer-institute-student-interns-achieve-dynamic-uas-model-using-nasa-supercomputer/
- **Small Scale, Big Effect**, *Ocean Bites*, February 28, 2018—Climate scientists at Cal Tech and NASA Jet Propulsion Laboratory used the Pleiades supercomputer to create global climate models of Earth's oceans on a submesoscale level.
  https://oceanbites.org/small-scale-big-effect/

HECC Utilization

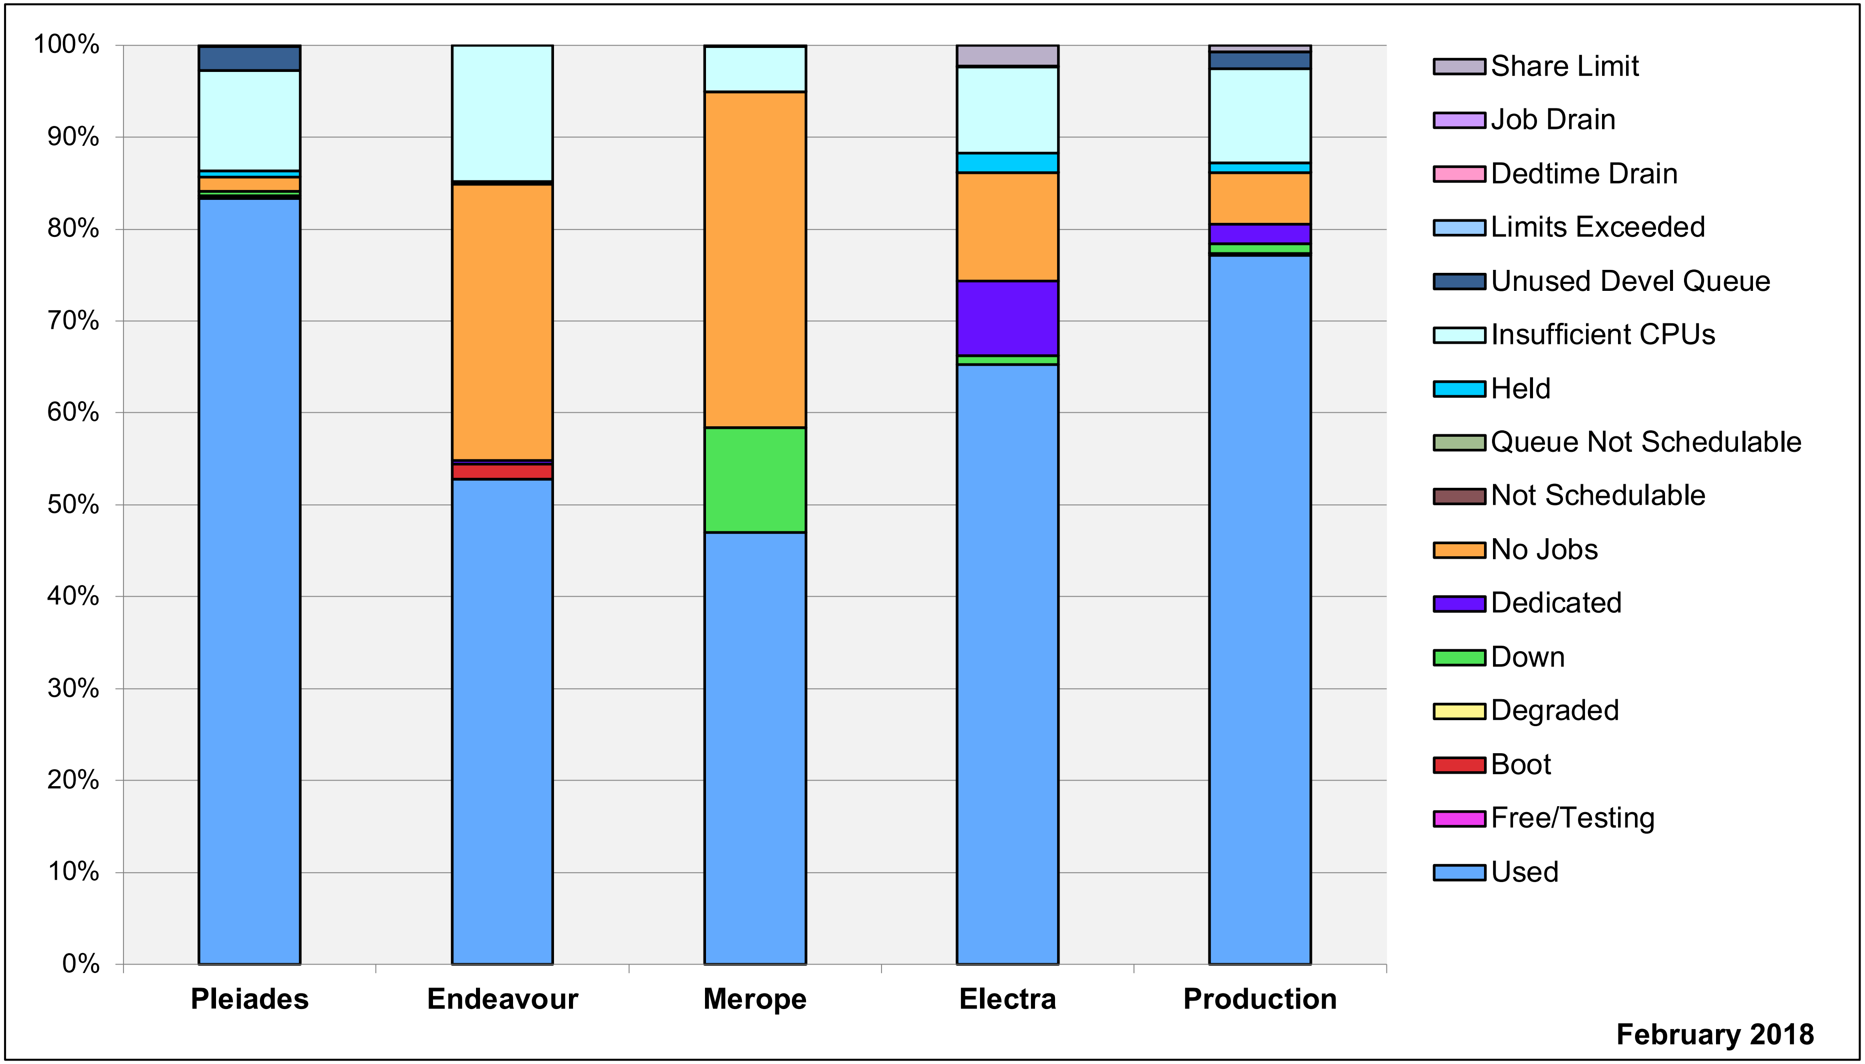

HECC Utilization Normalized to 30-Day Month

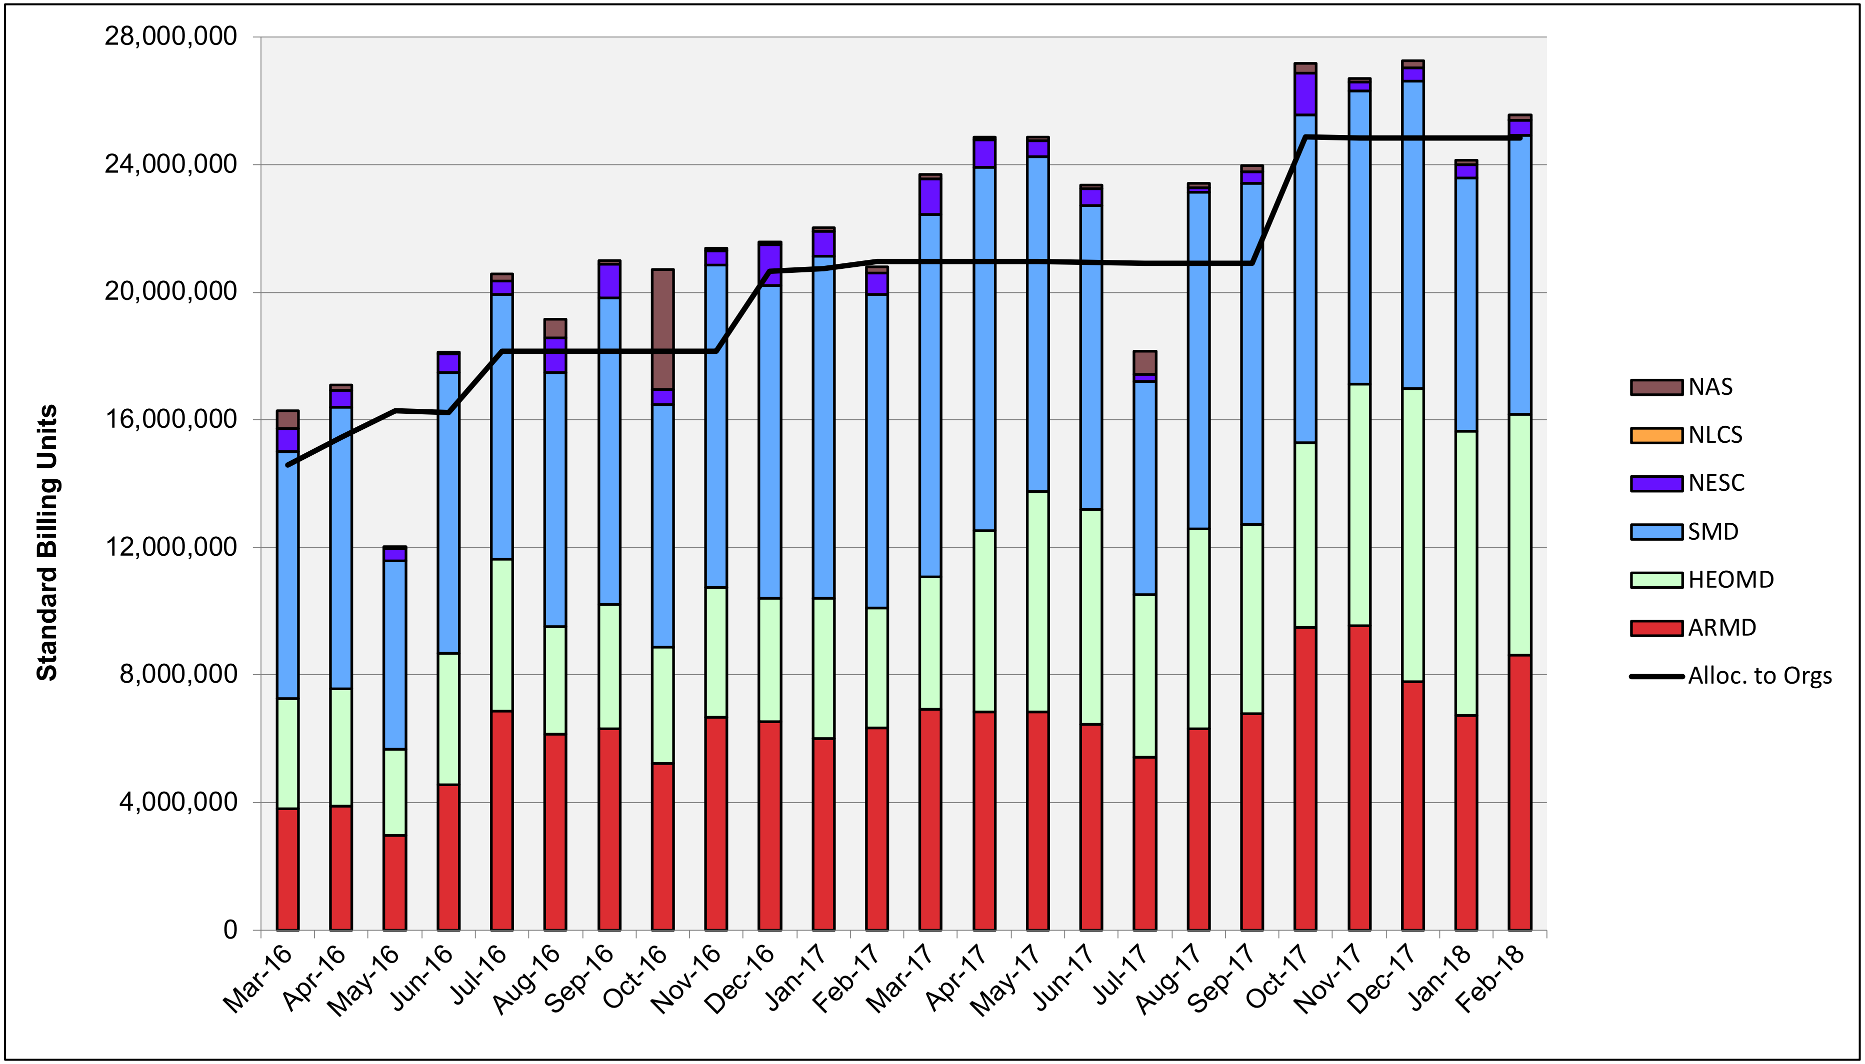

HECC Utilization Normalized to 30-Day Month

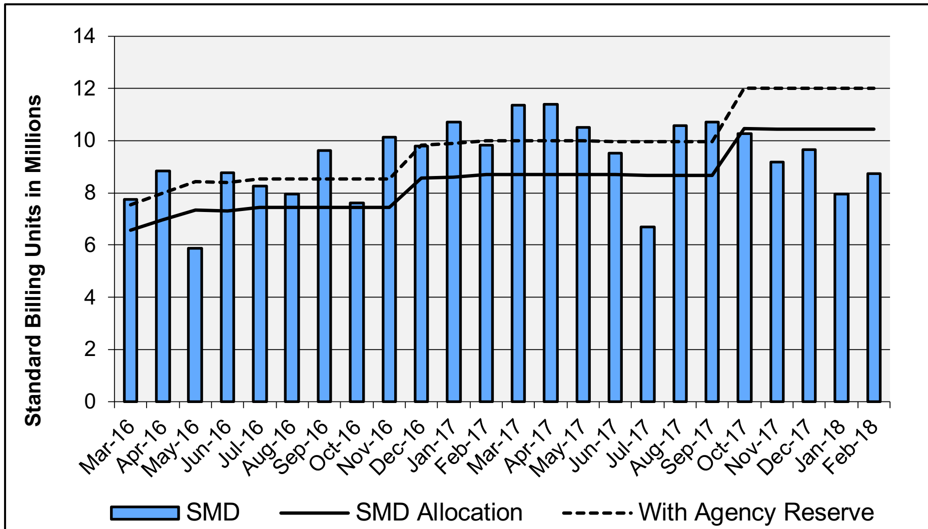

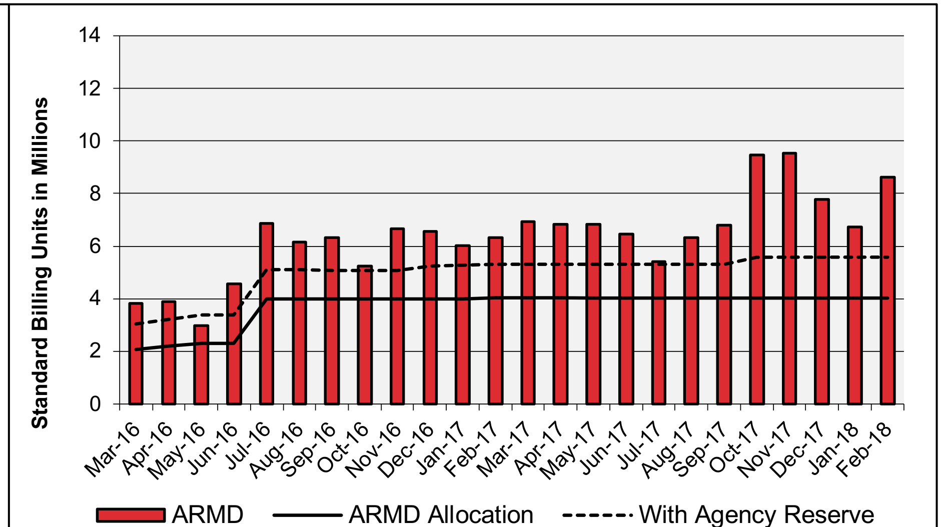

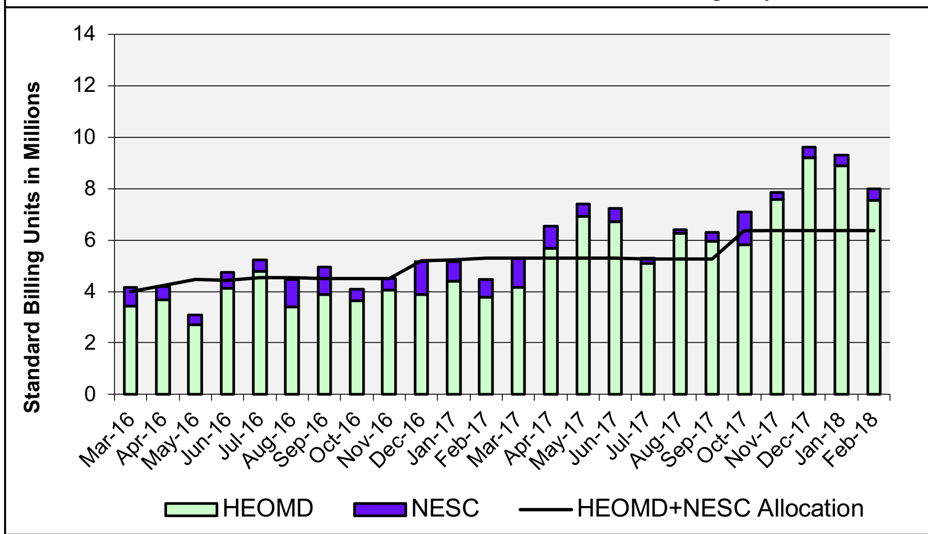

Tape Archive Status

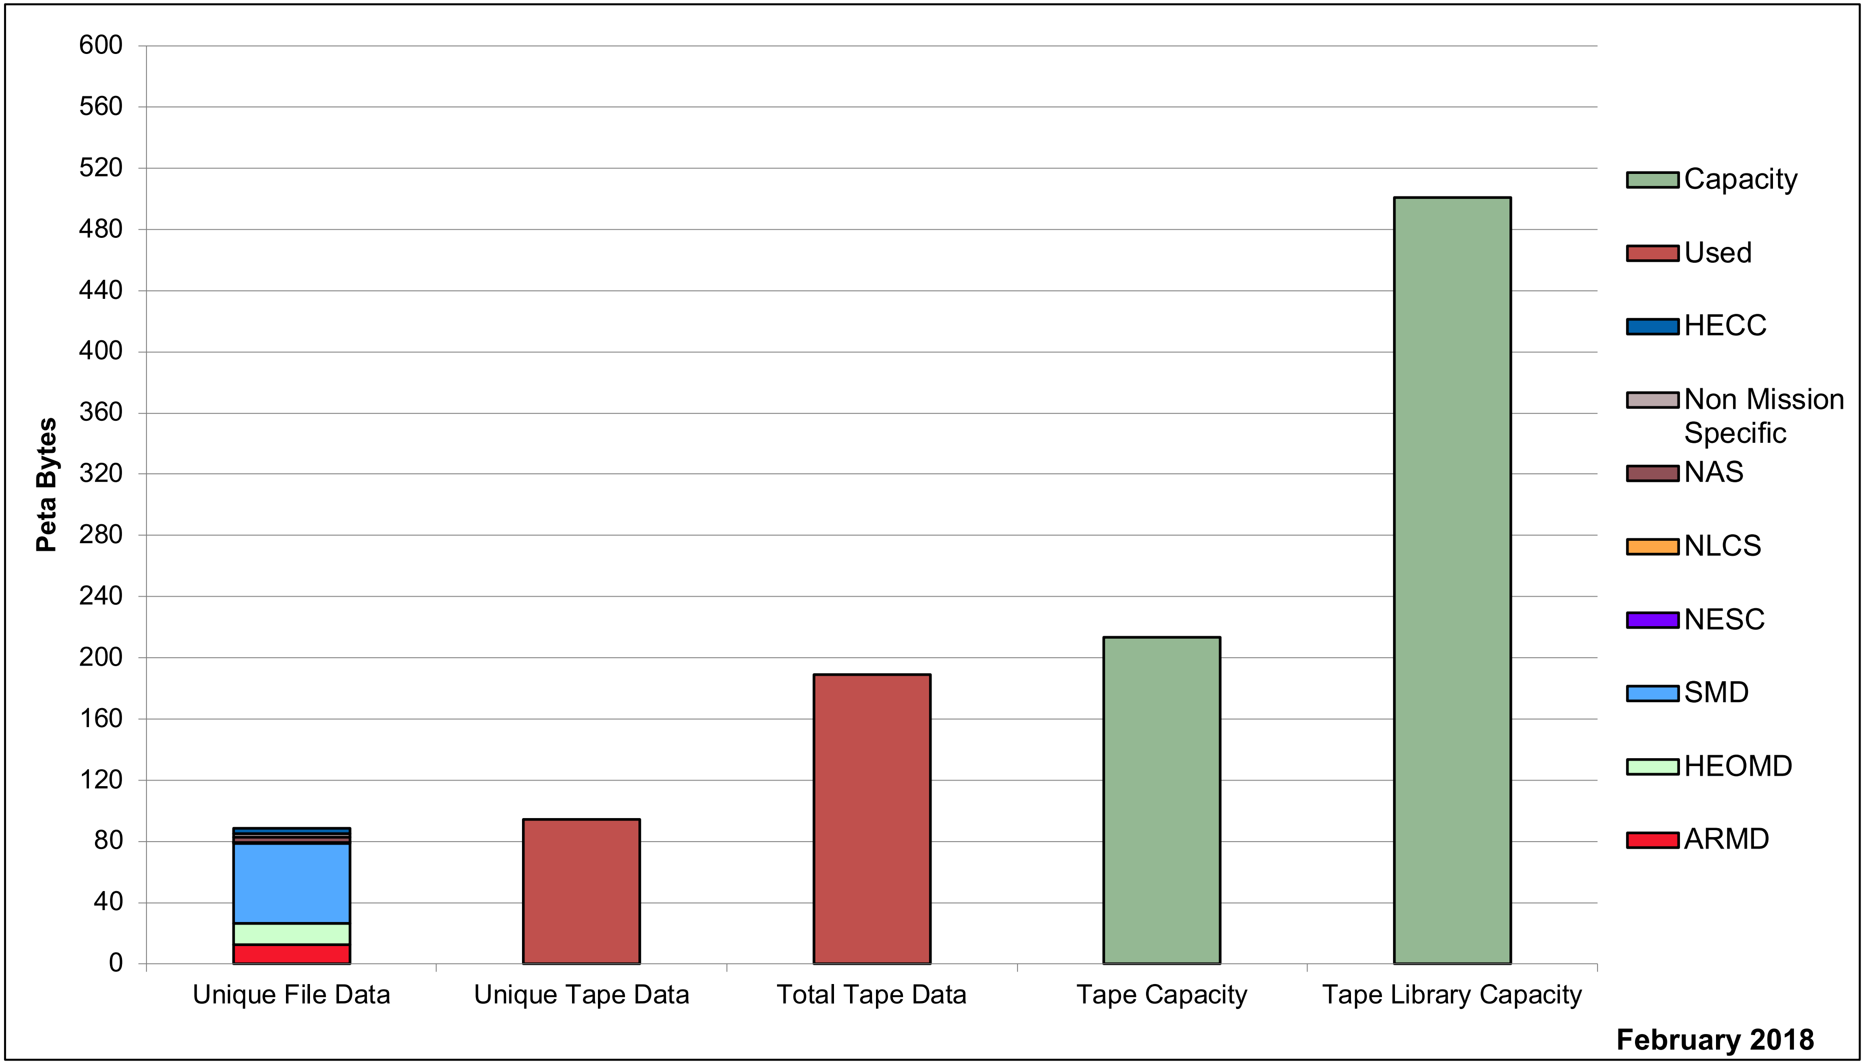

Tape Archive Status

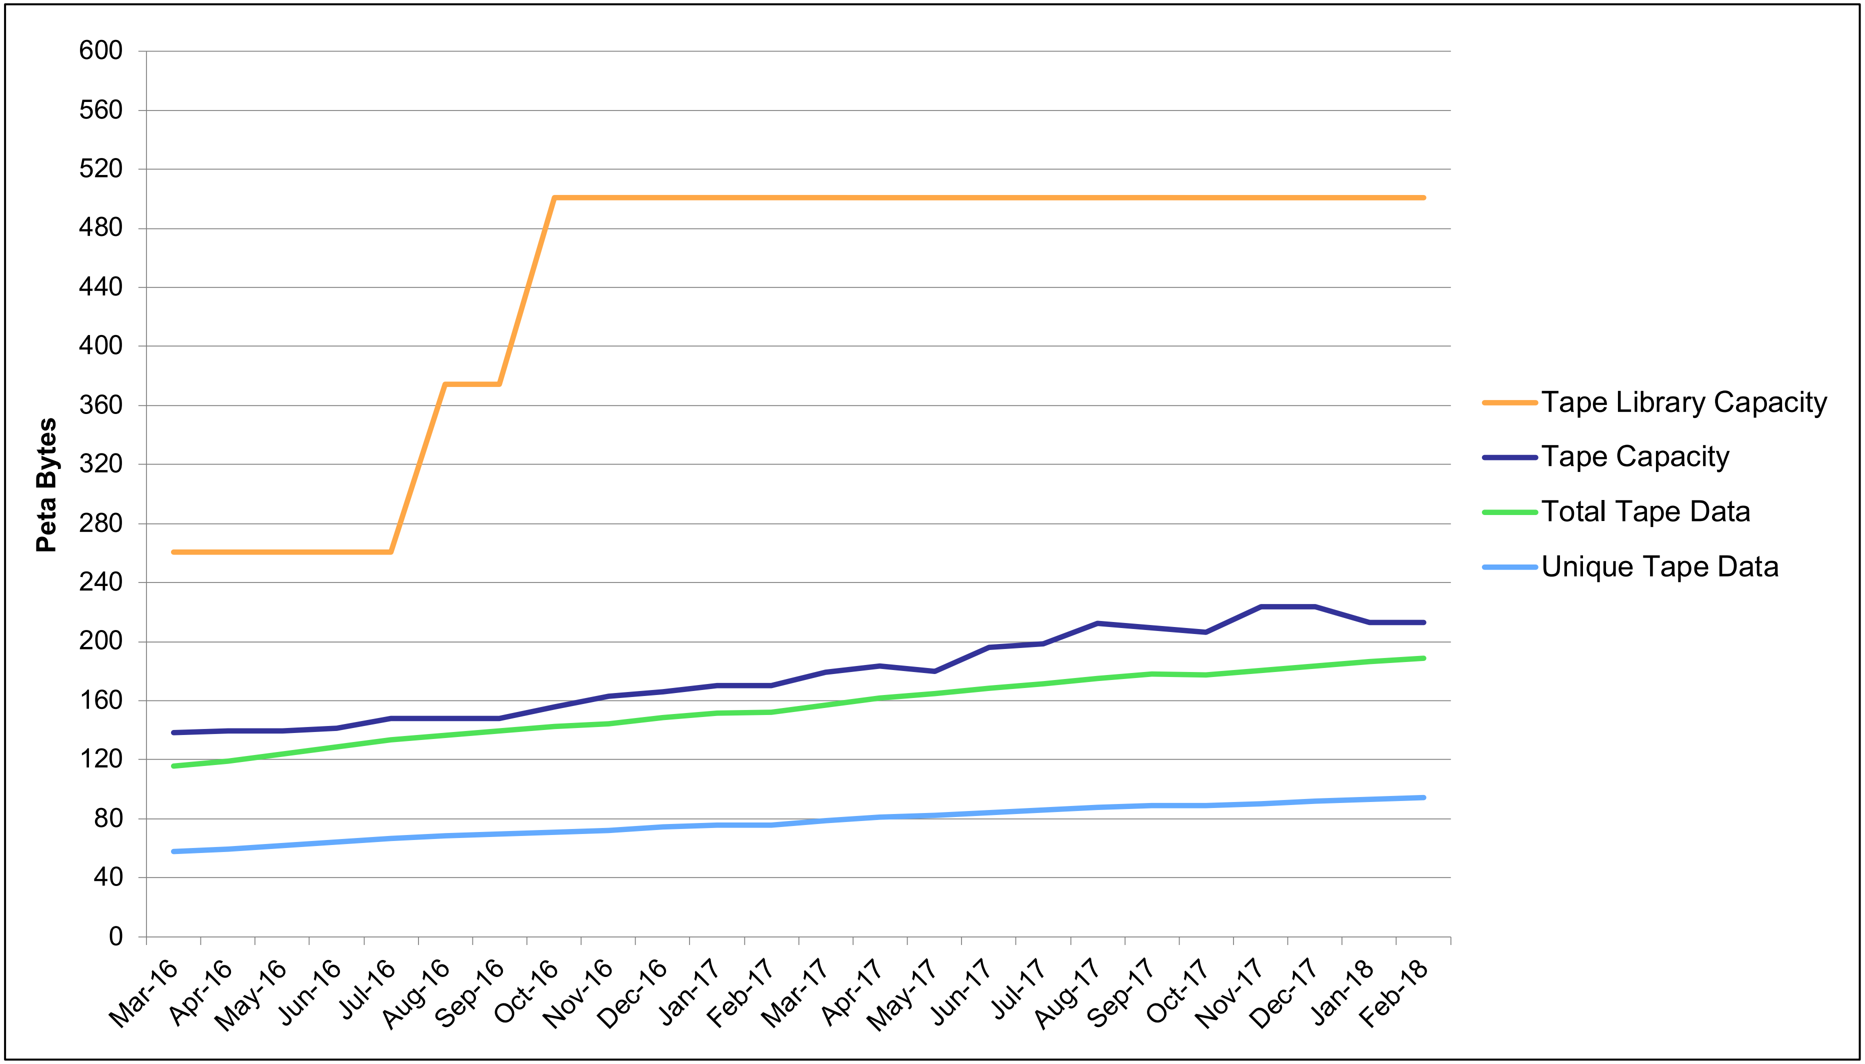

Pleiades:
SBUs Reported, Normalized to 30-Day Month

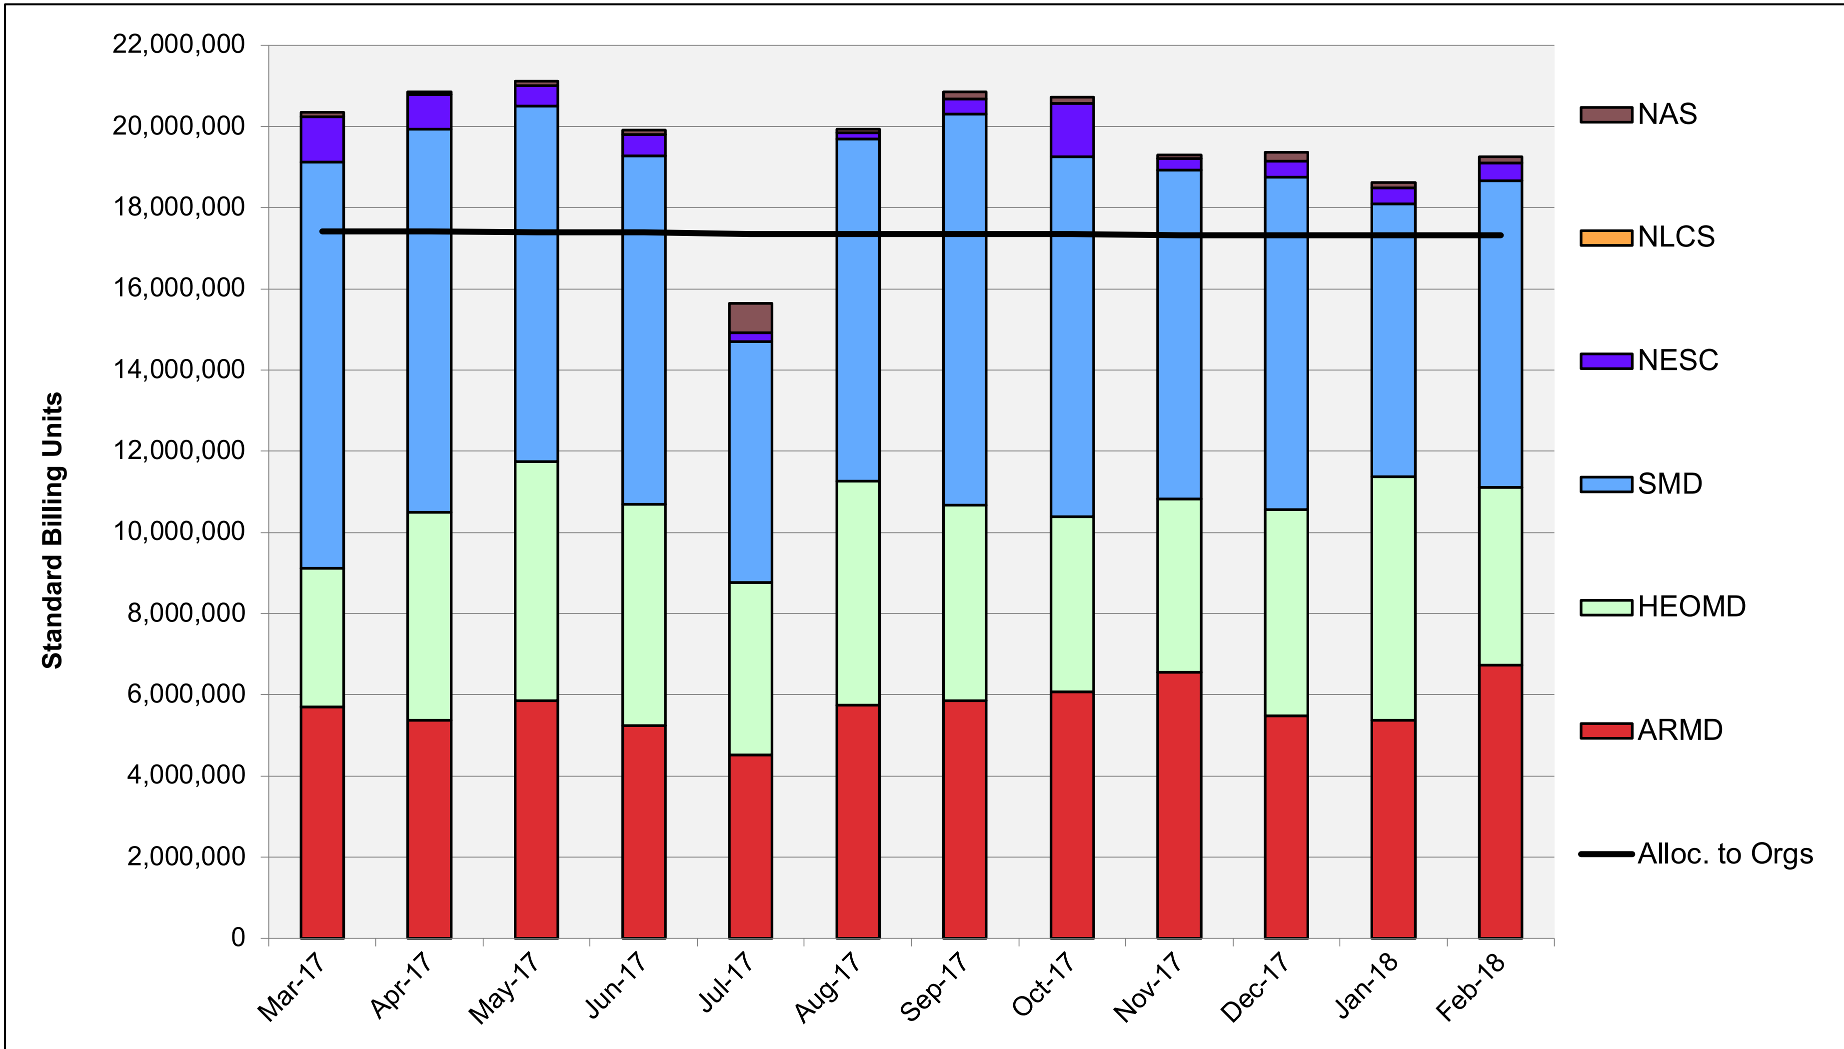

Pleiades: Devel Queue Utilization

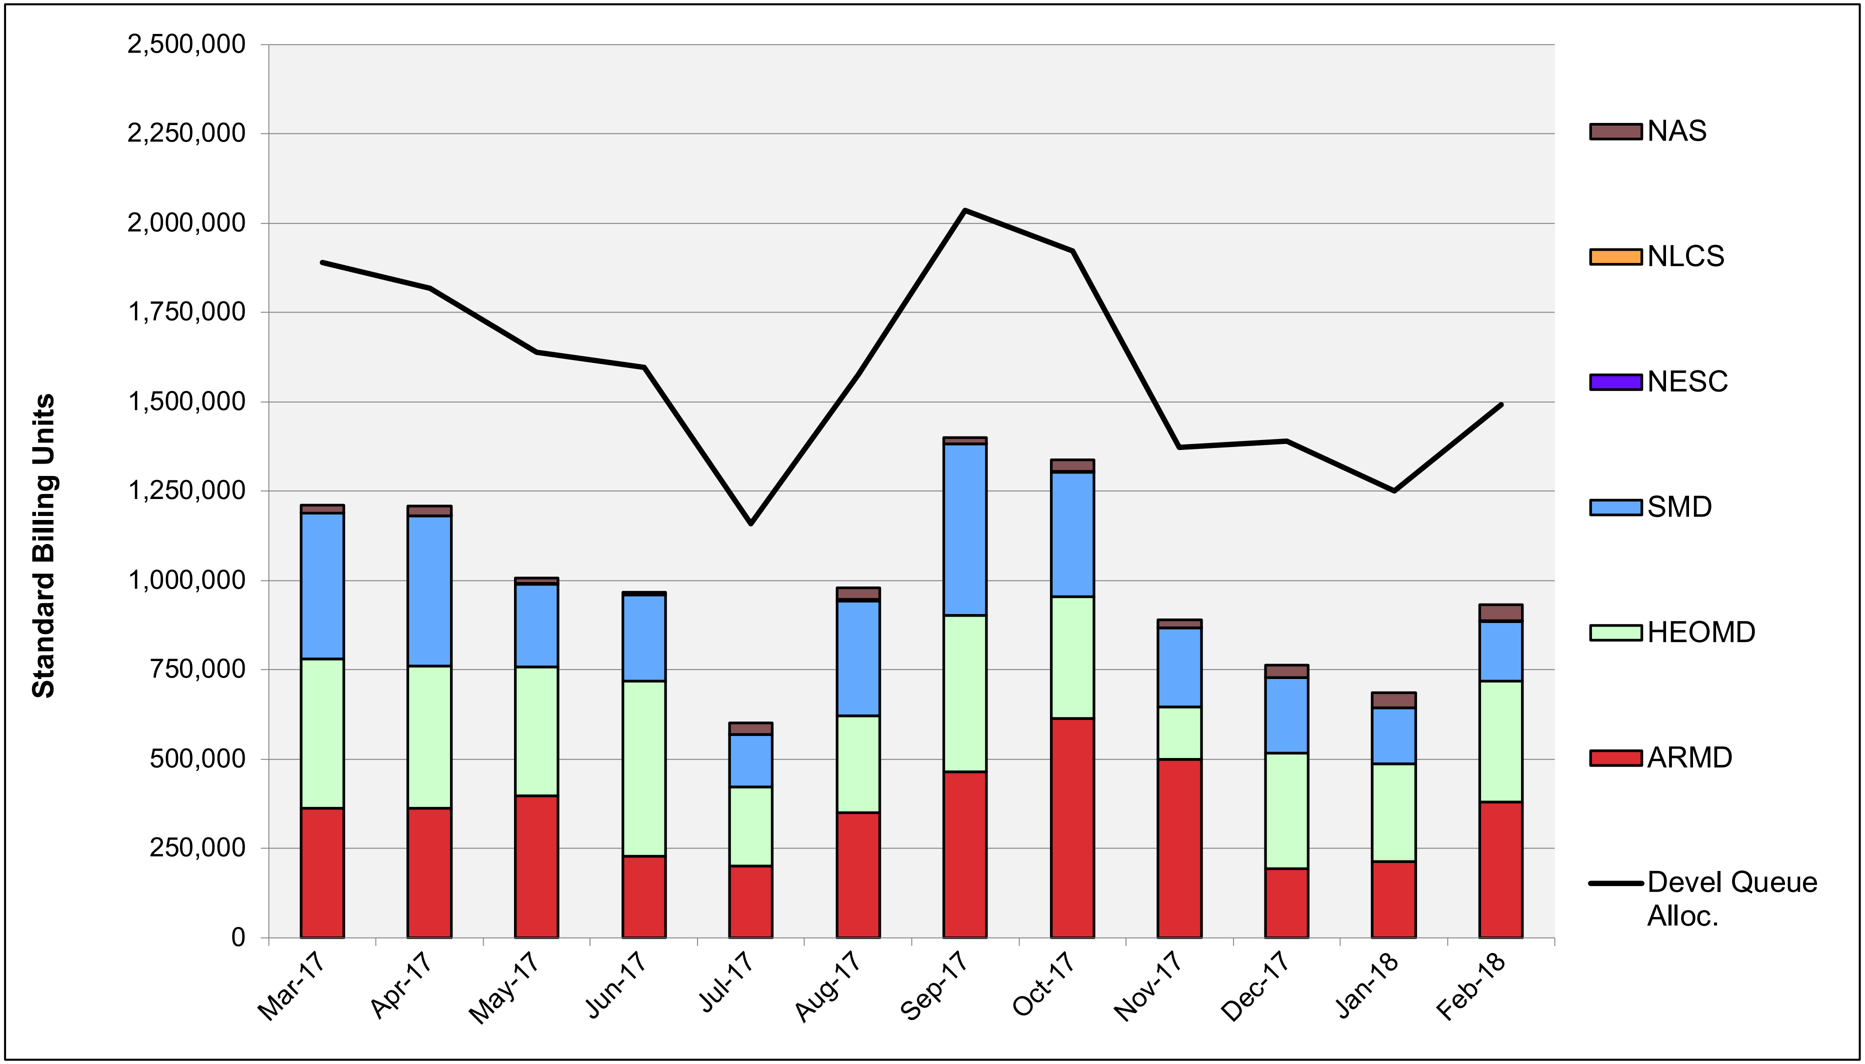

Pleiades: Monthly Utilization by Job Length

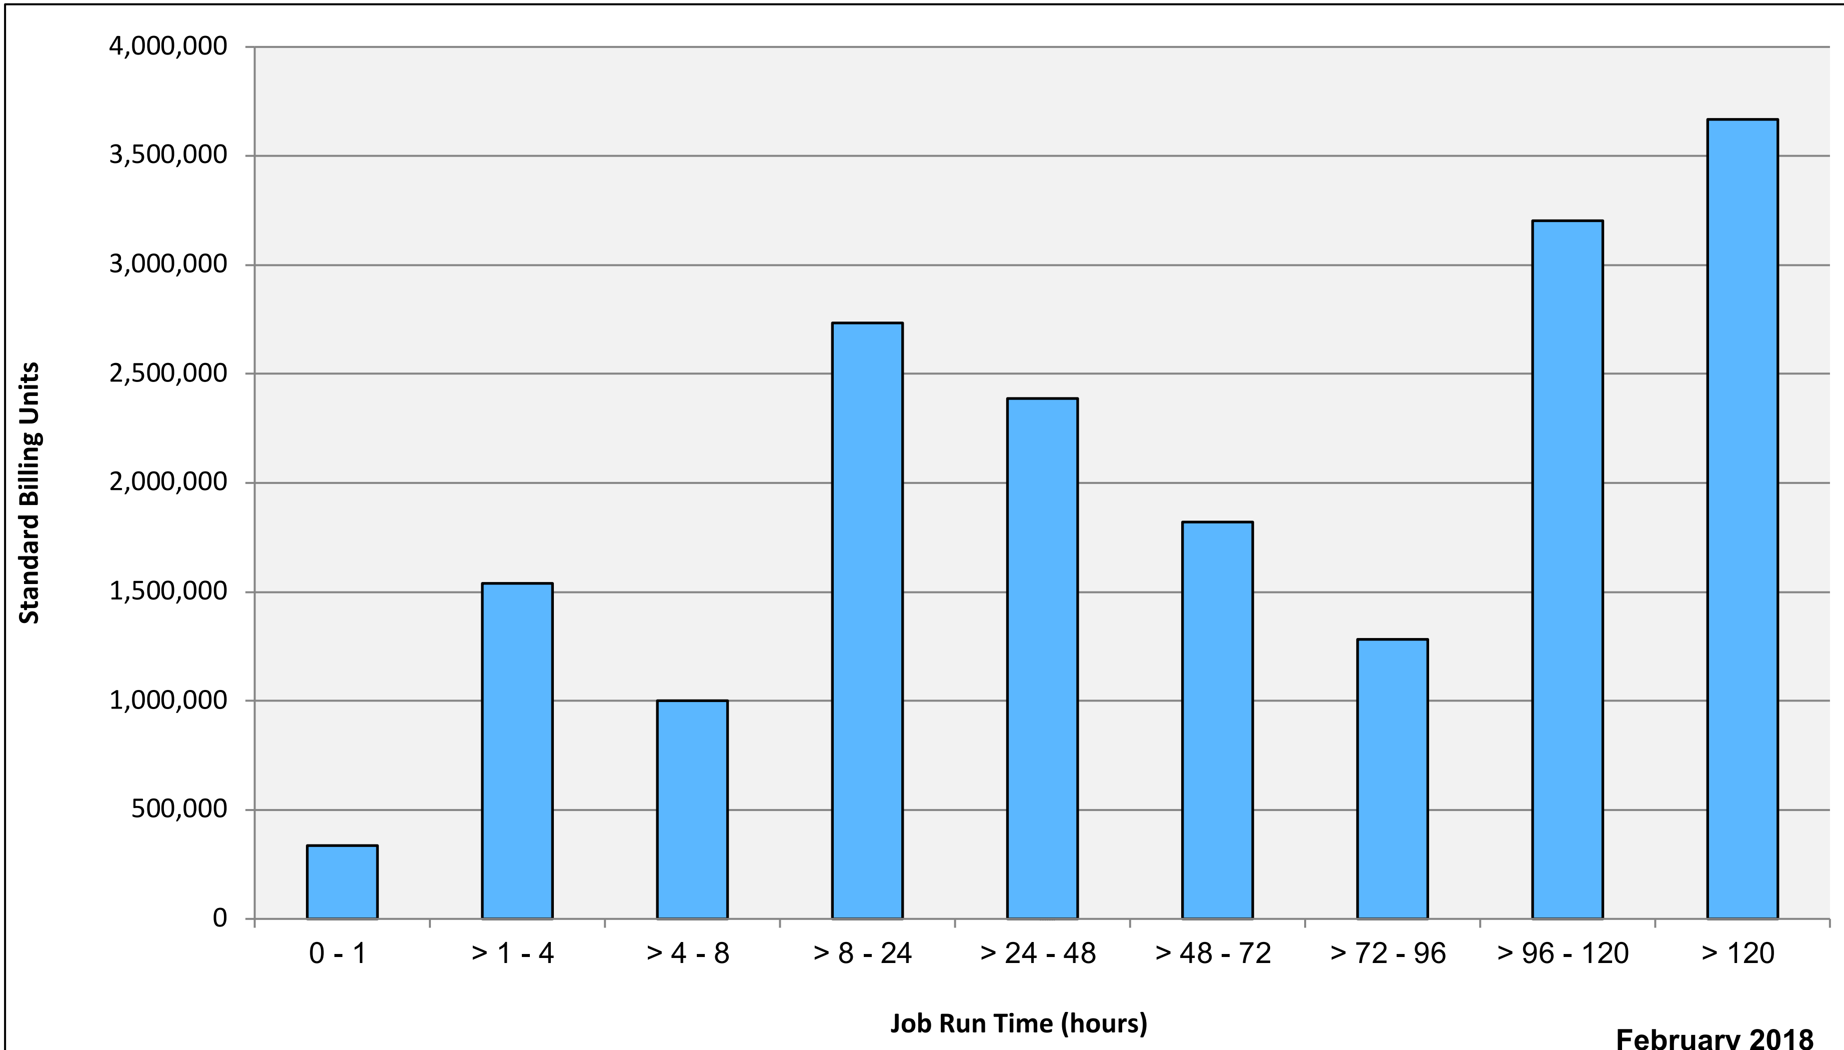

Pleiades: Monthly Utilization by Size and Mission

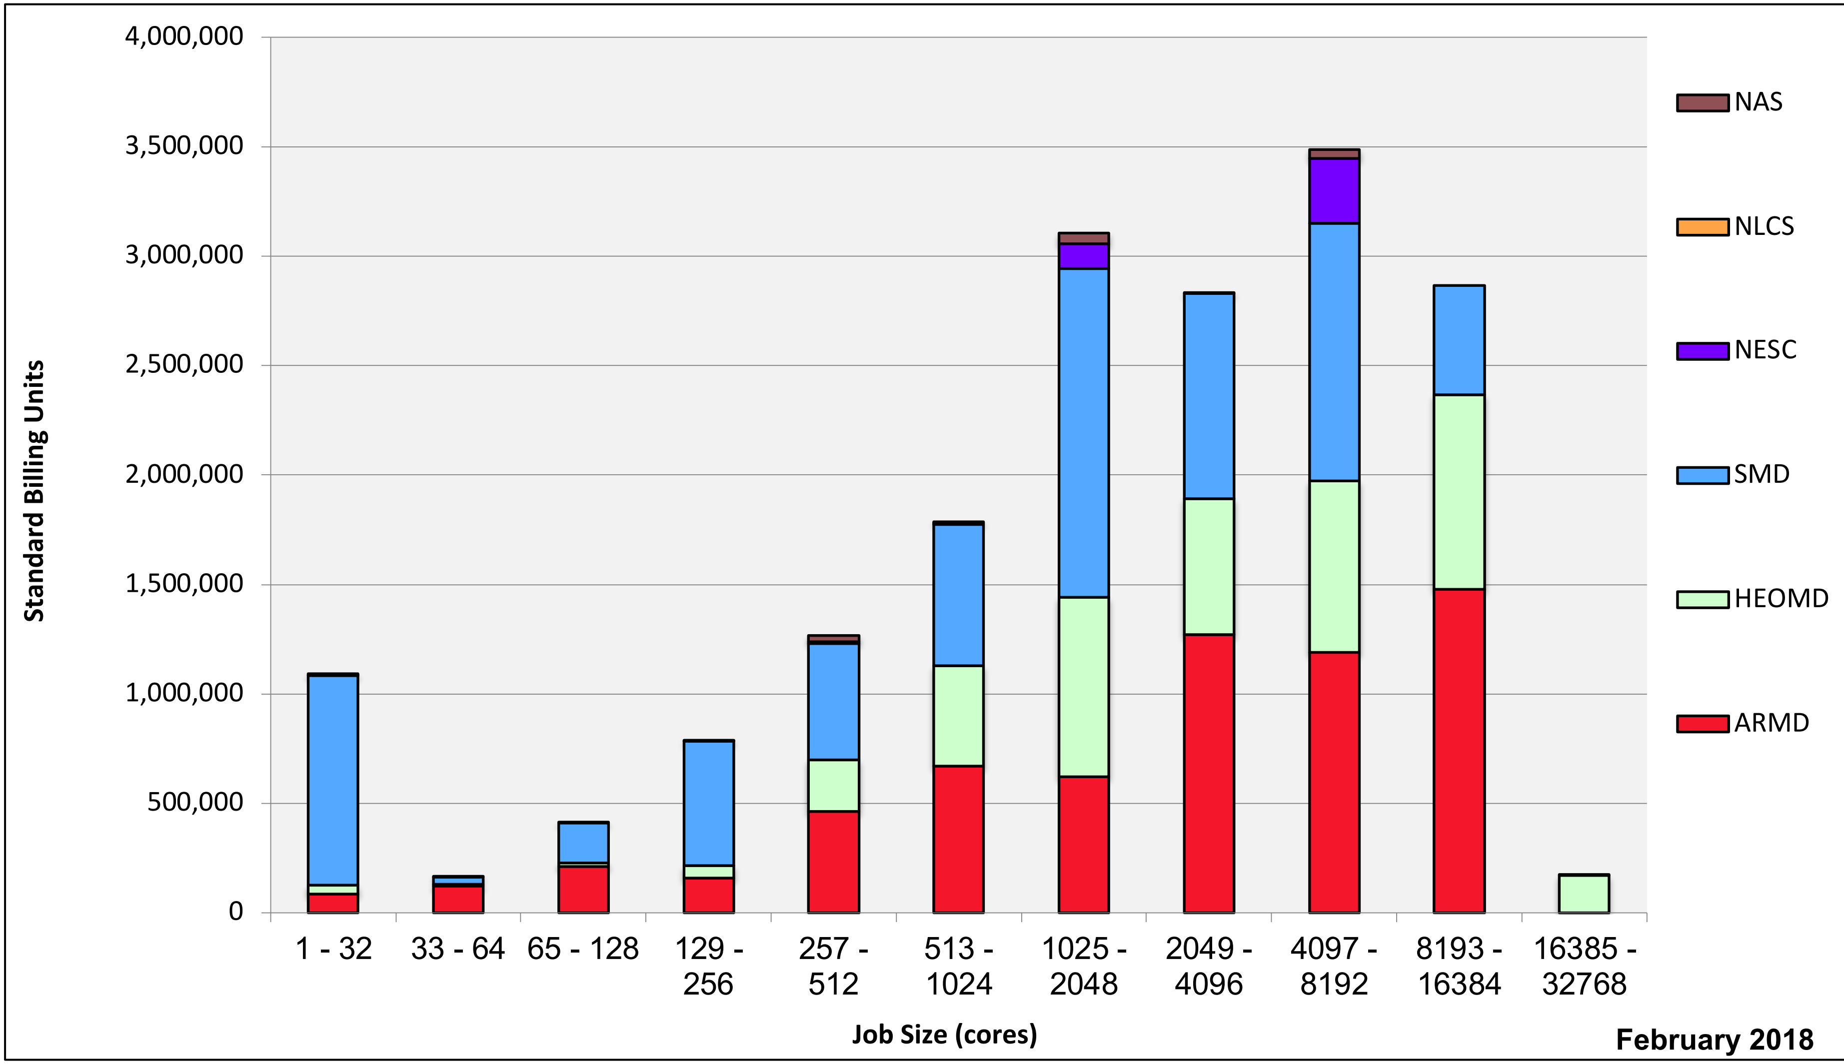

Pleiades: Monthly Utilization by Size and Length

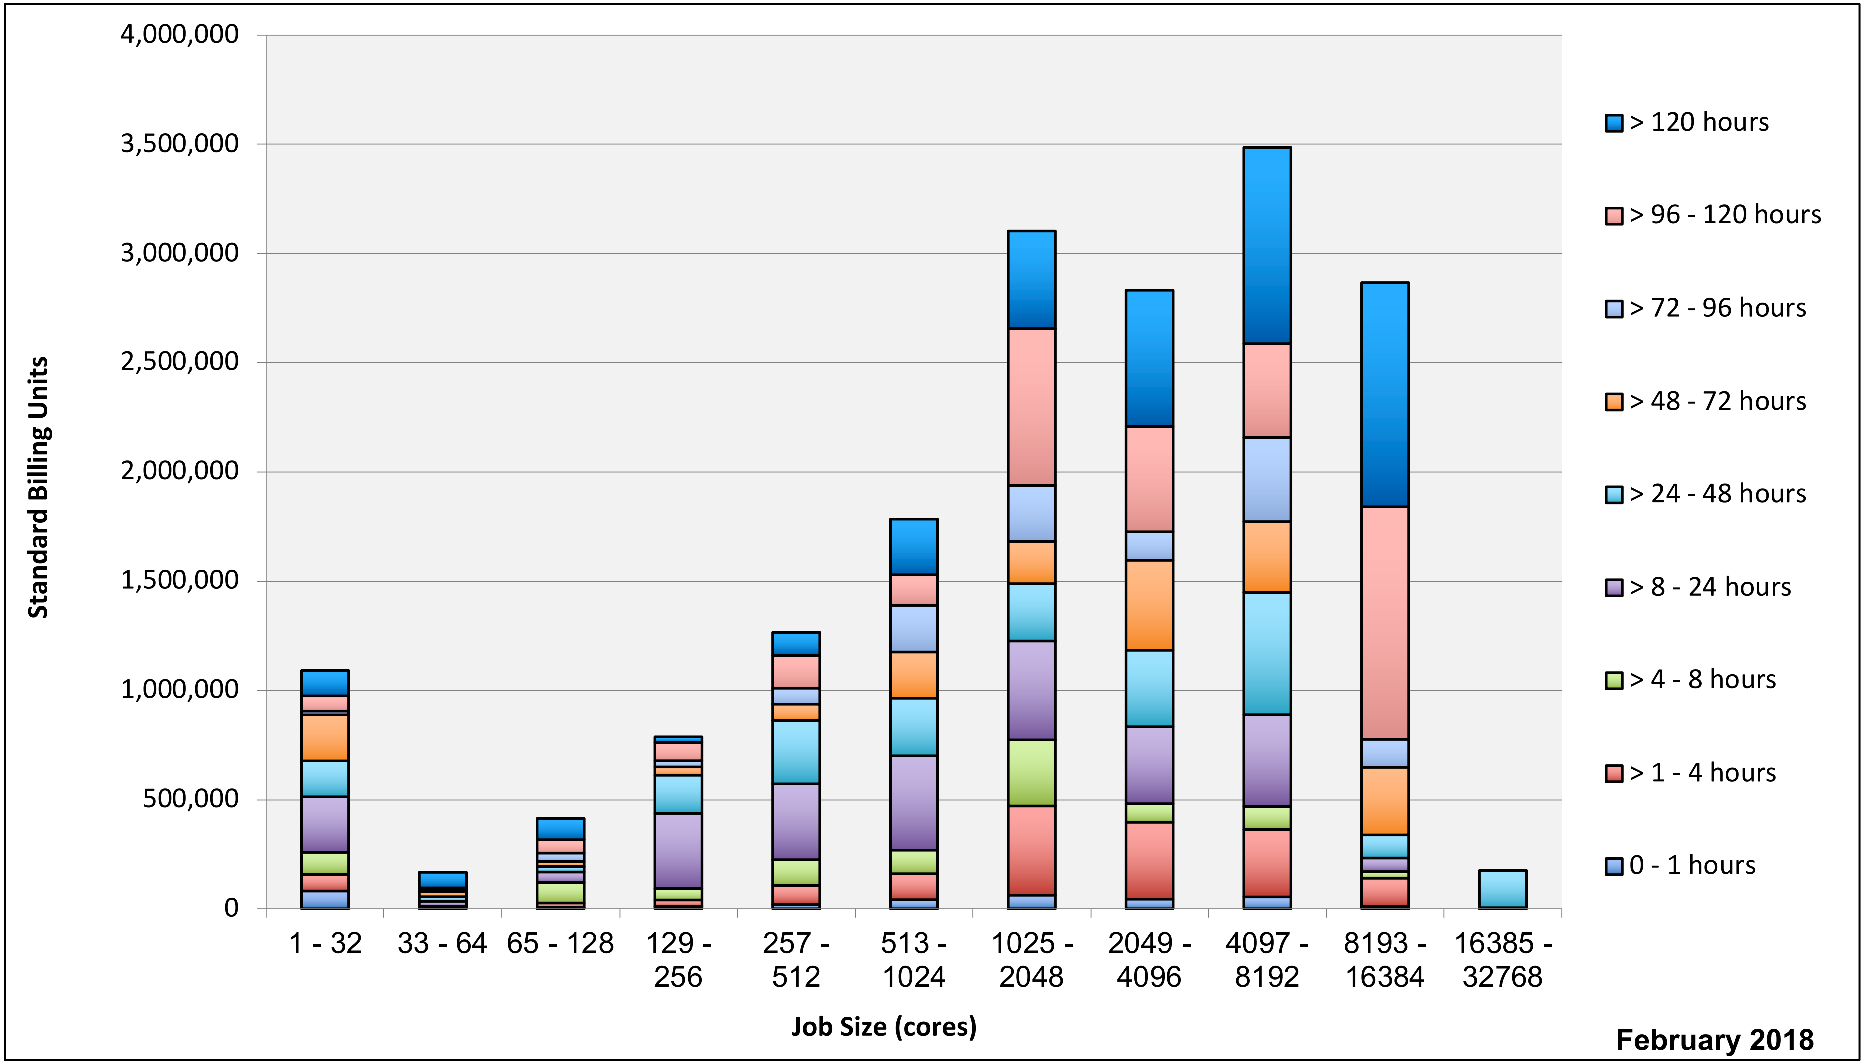

Pleiades:
Average Time to Clear All Jobs

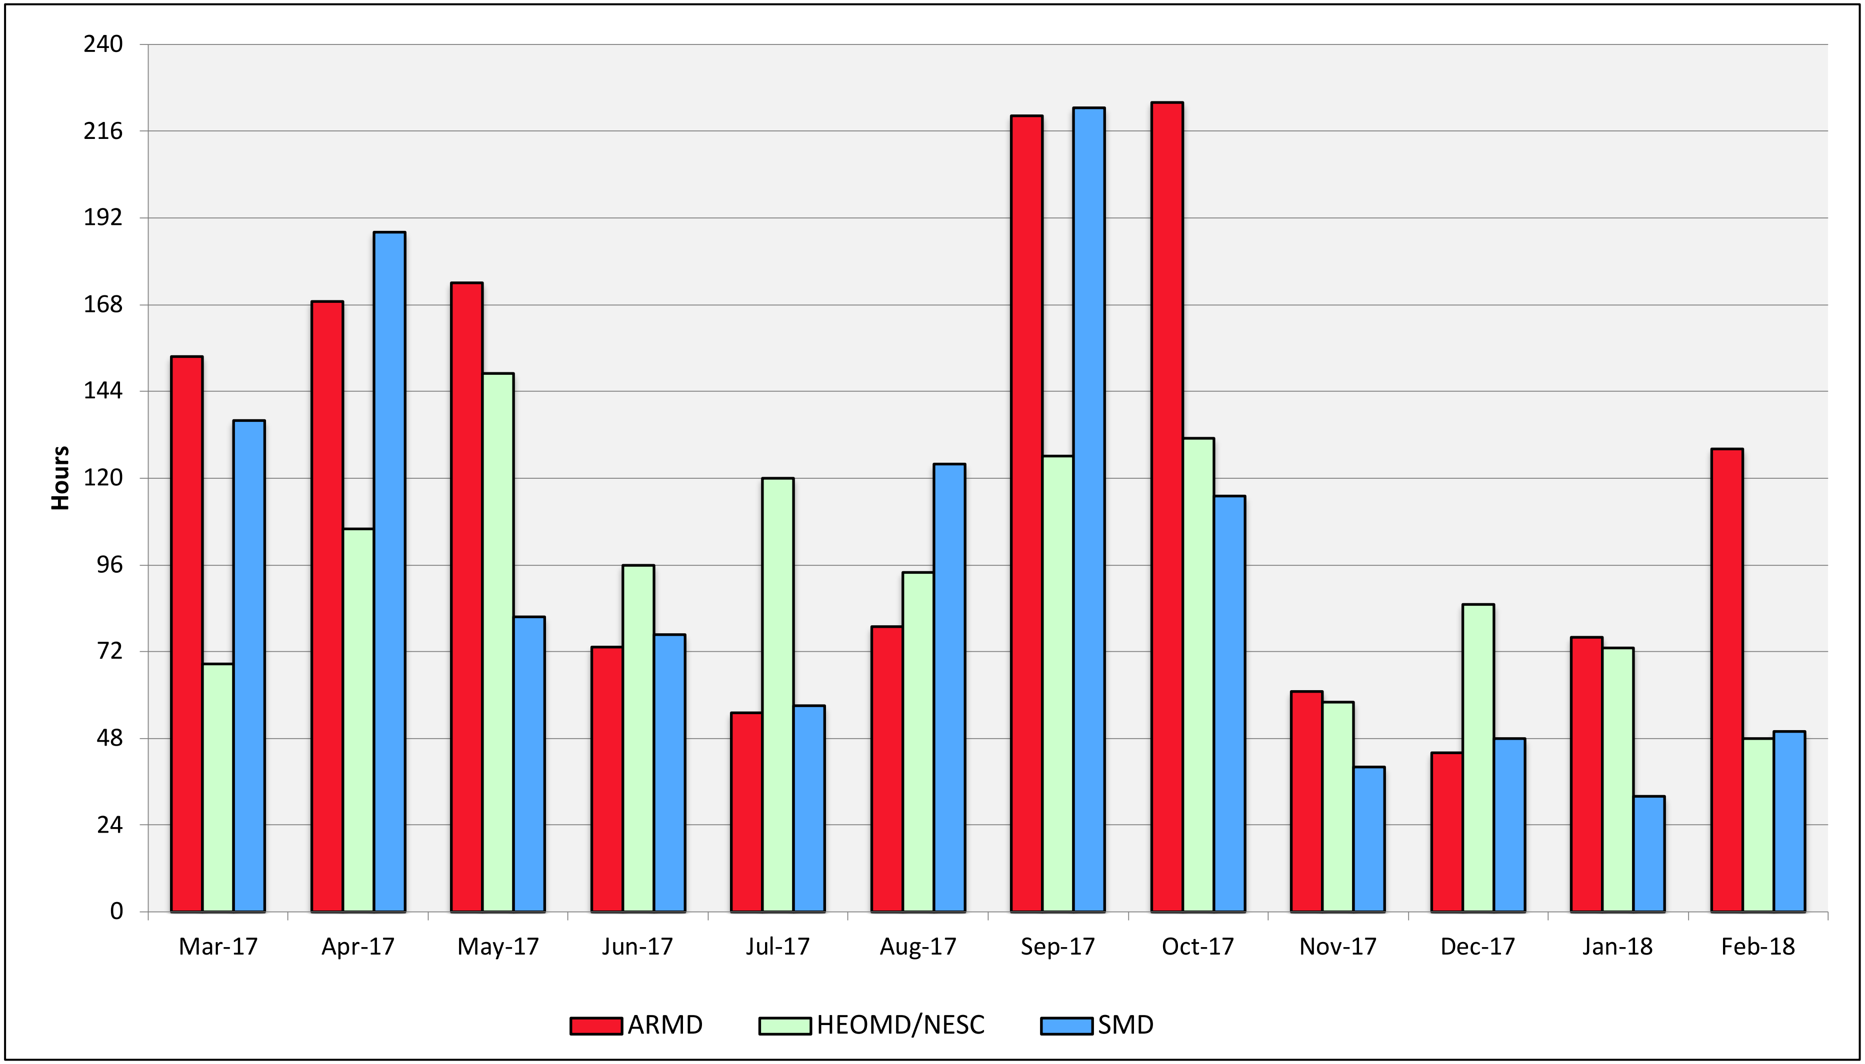

Pleiades: Average Expansion Factor

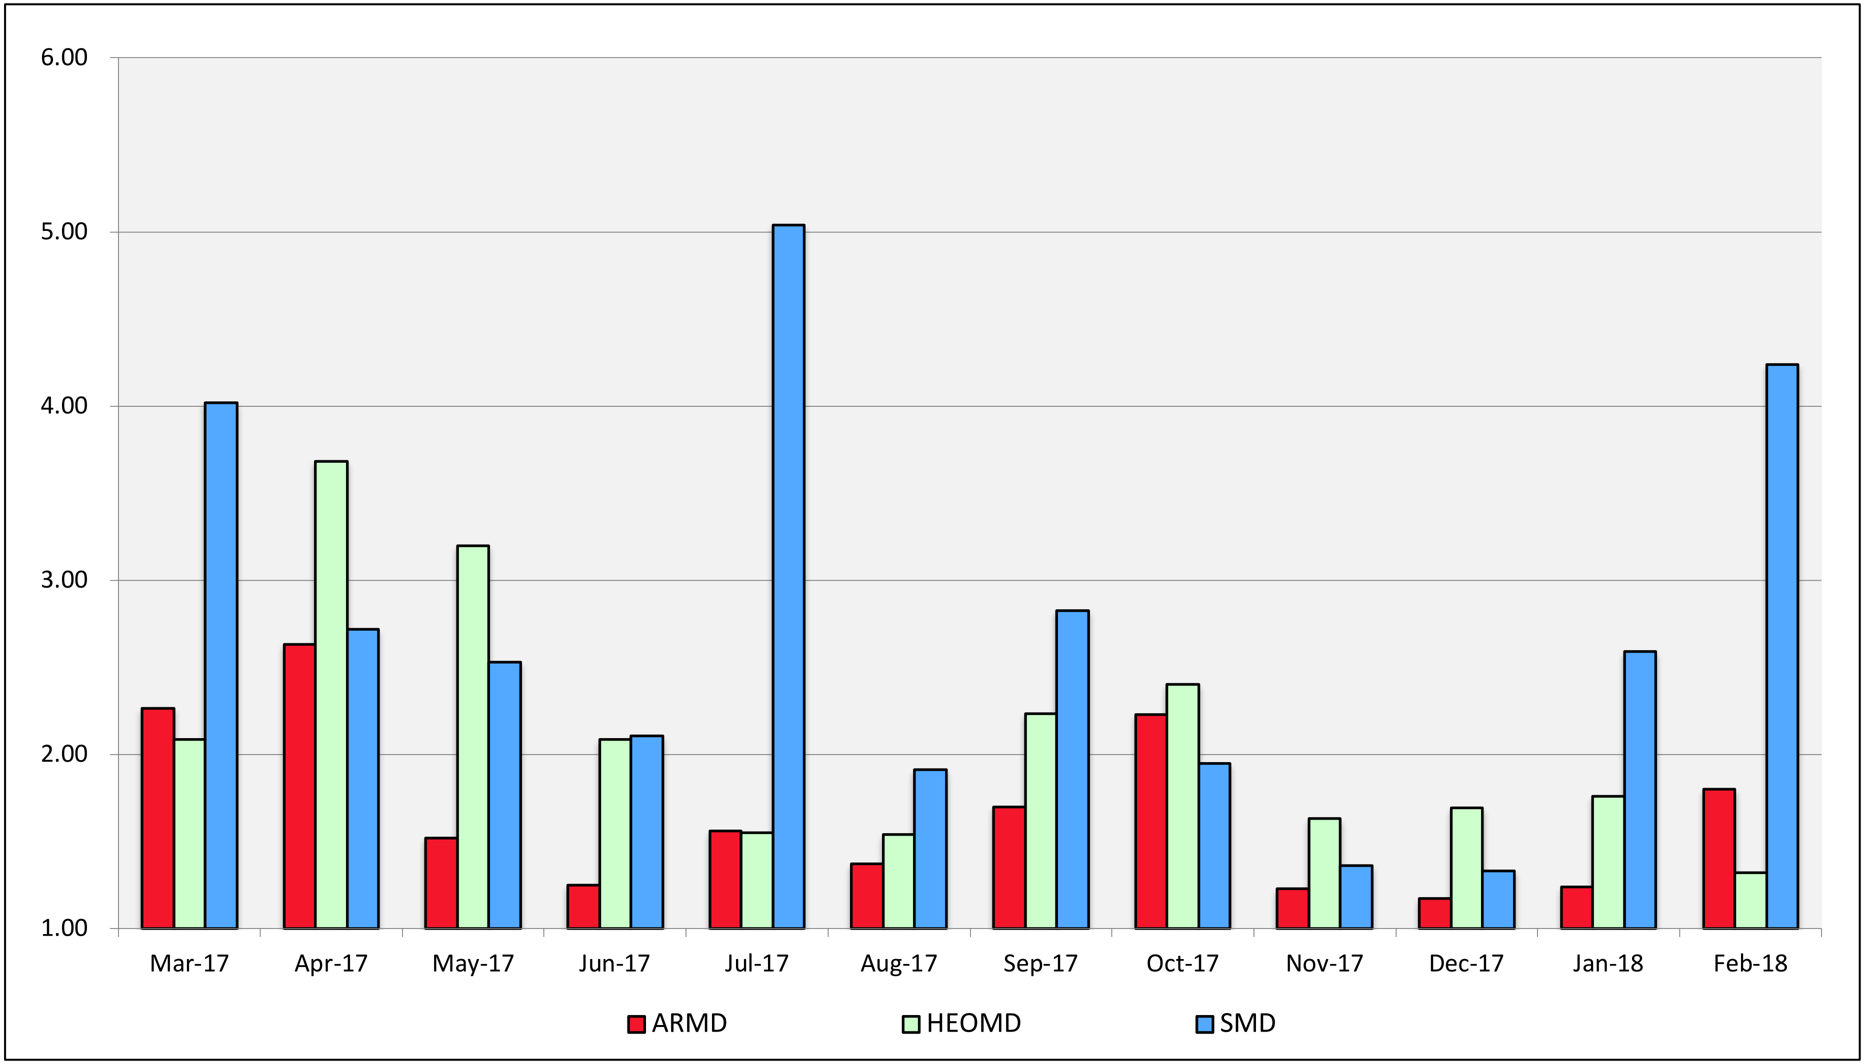

Electra:
SBUs Reported, Normalized to 30-Day Month

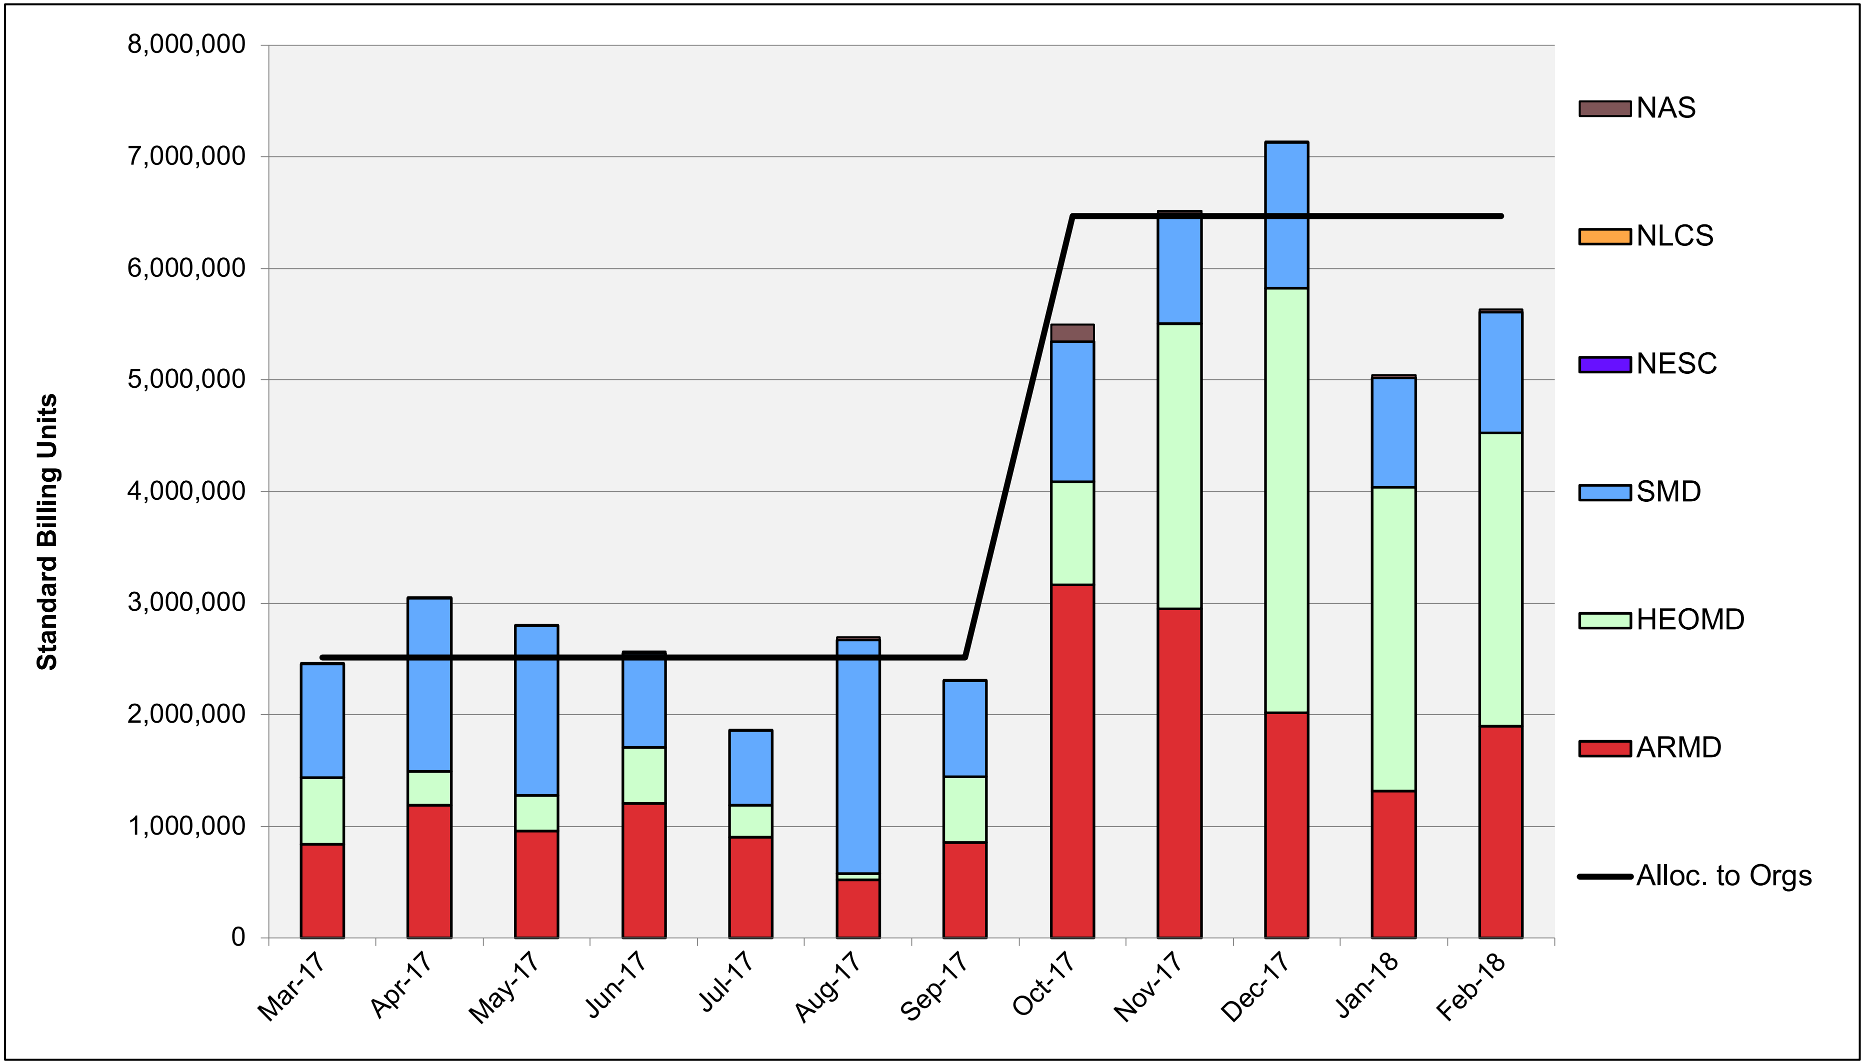

Electra:
Devel Queue Utilization

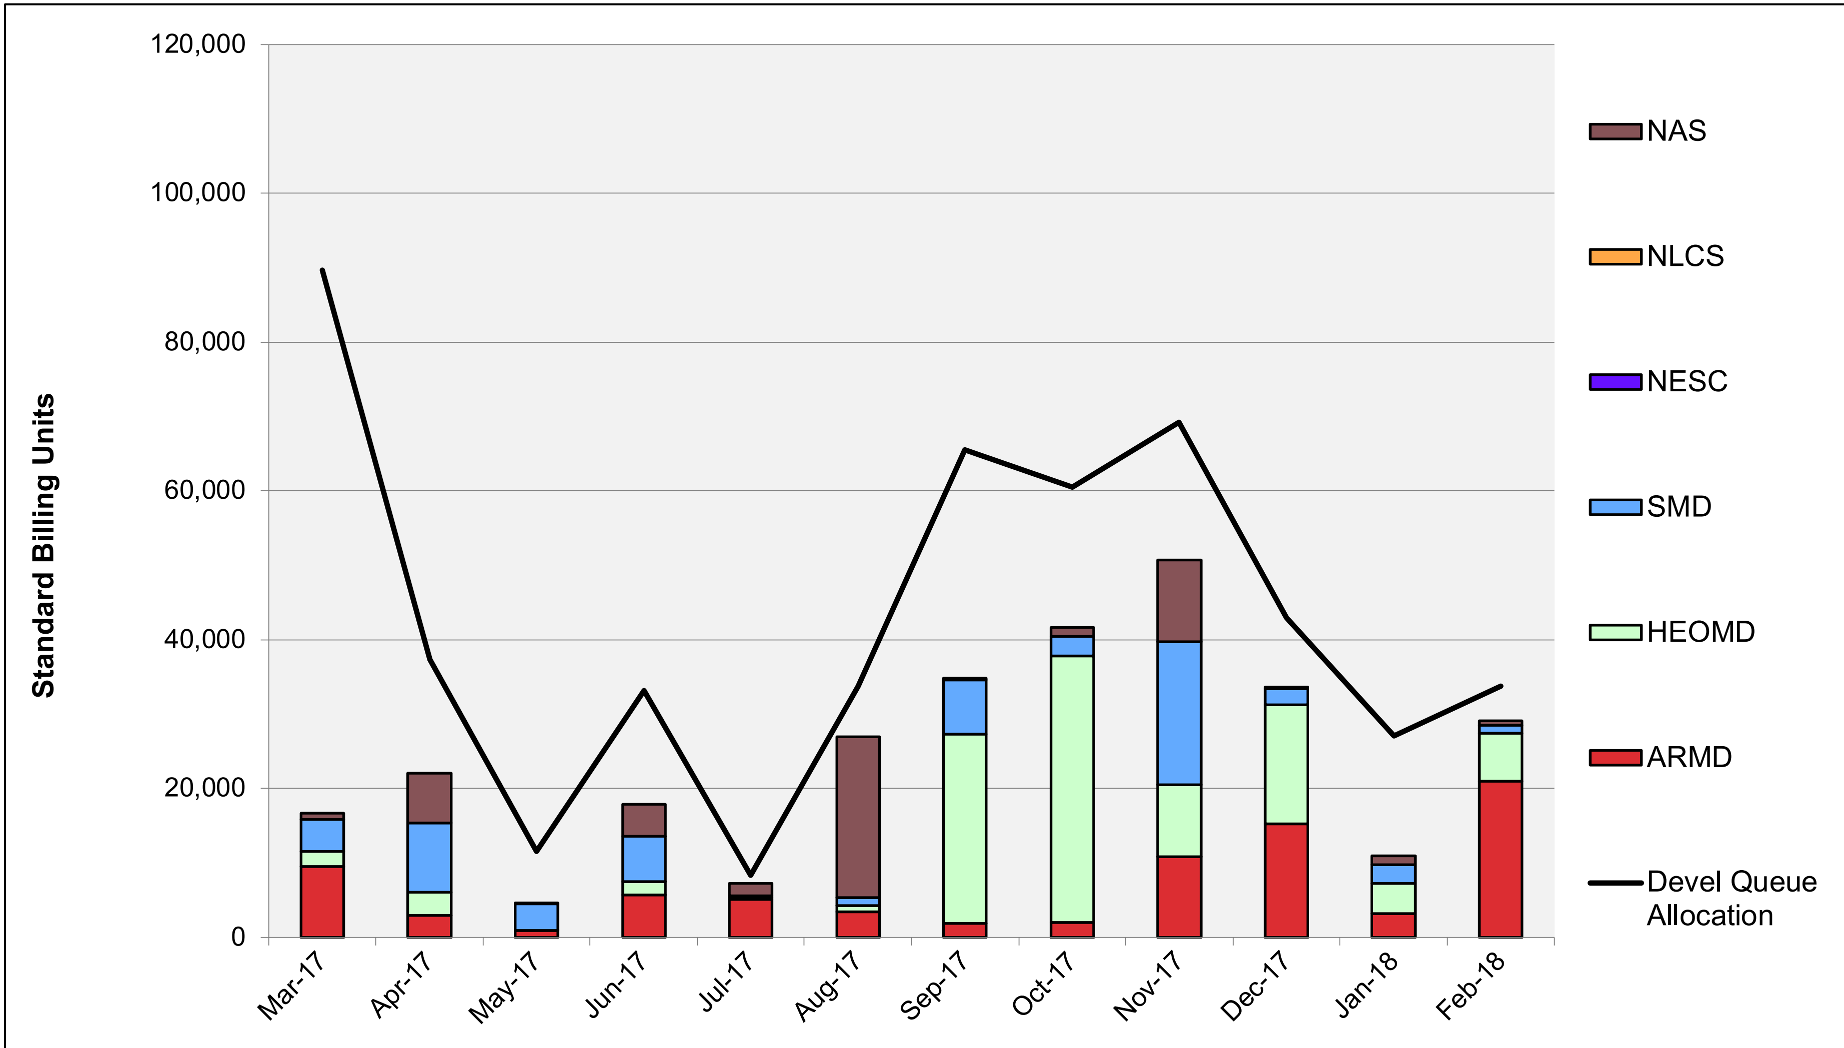

Electra: Monthly Utilization by Job Length

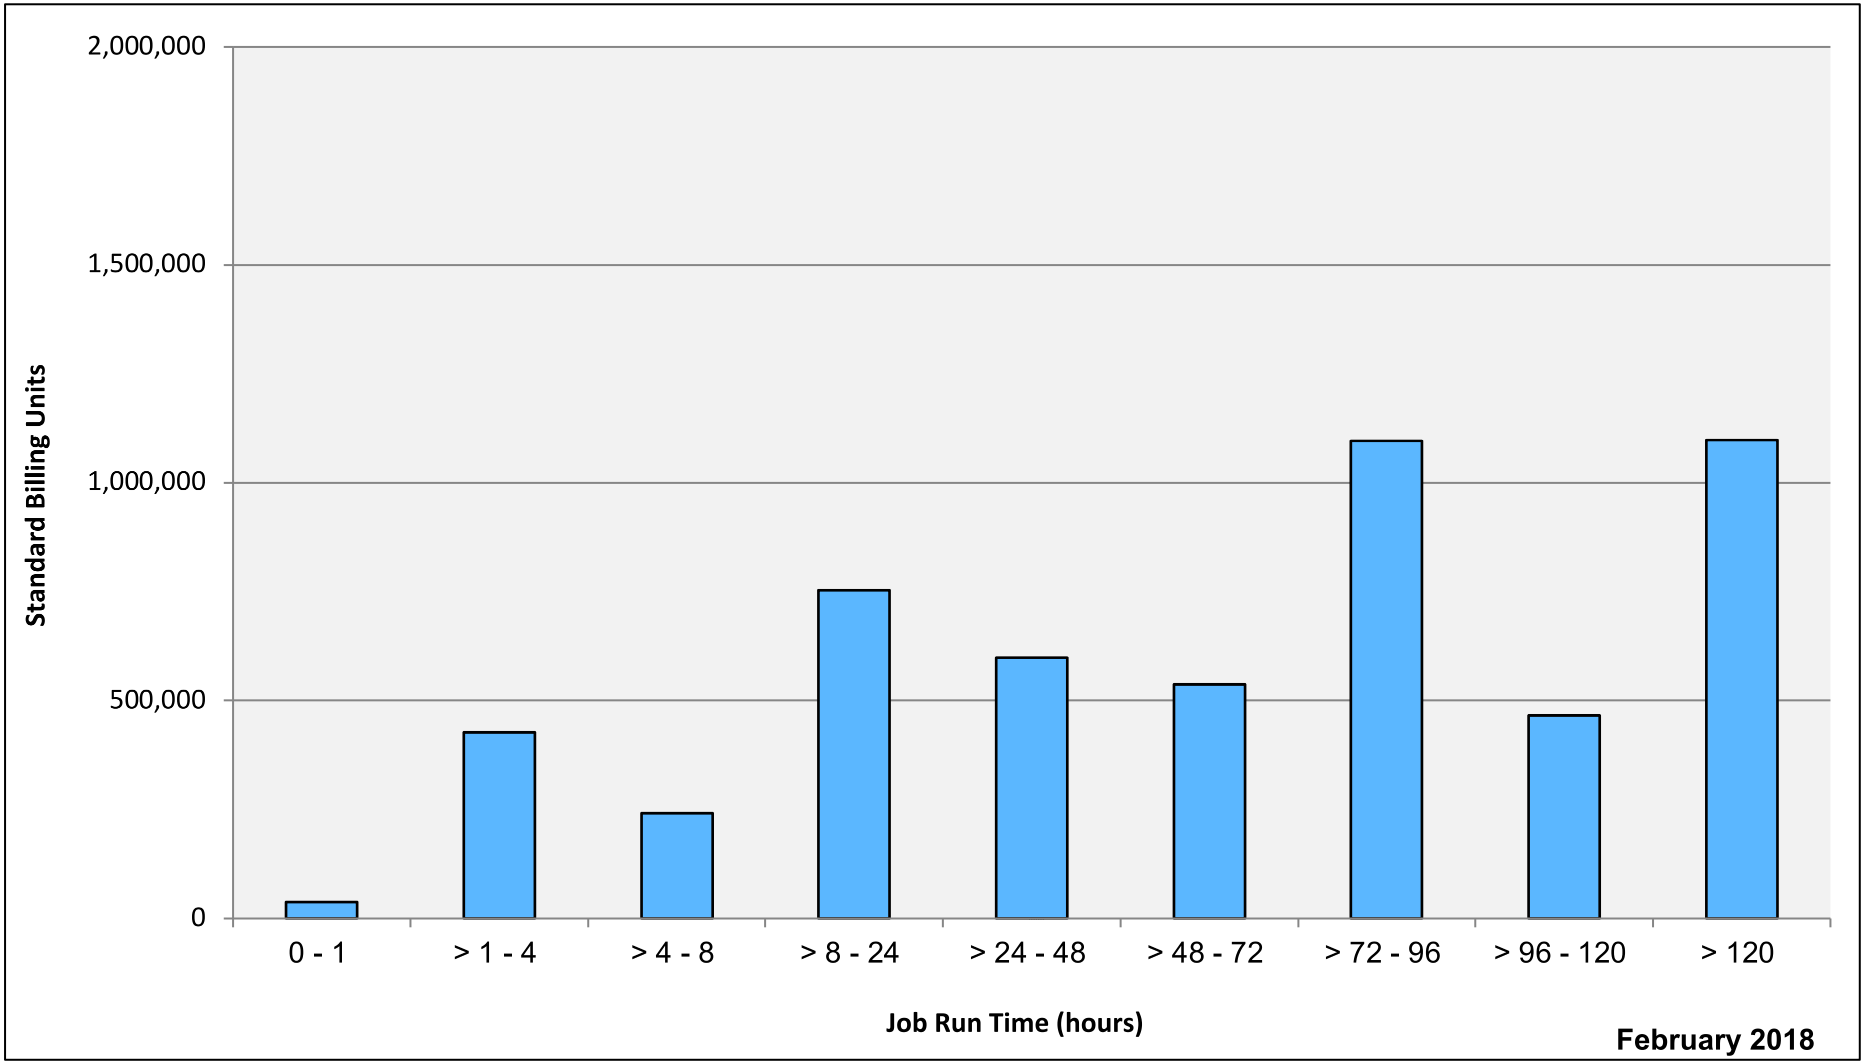

Electra: Monthly Utilization by Size and Mission

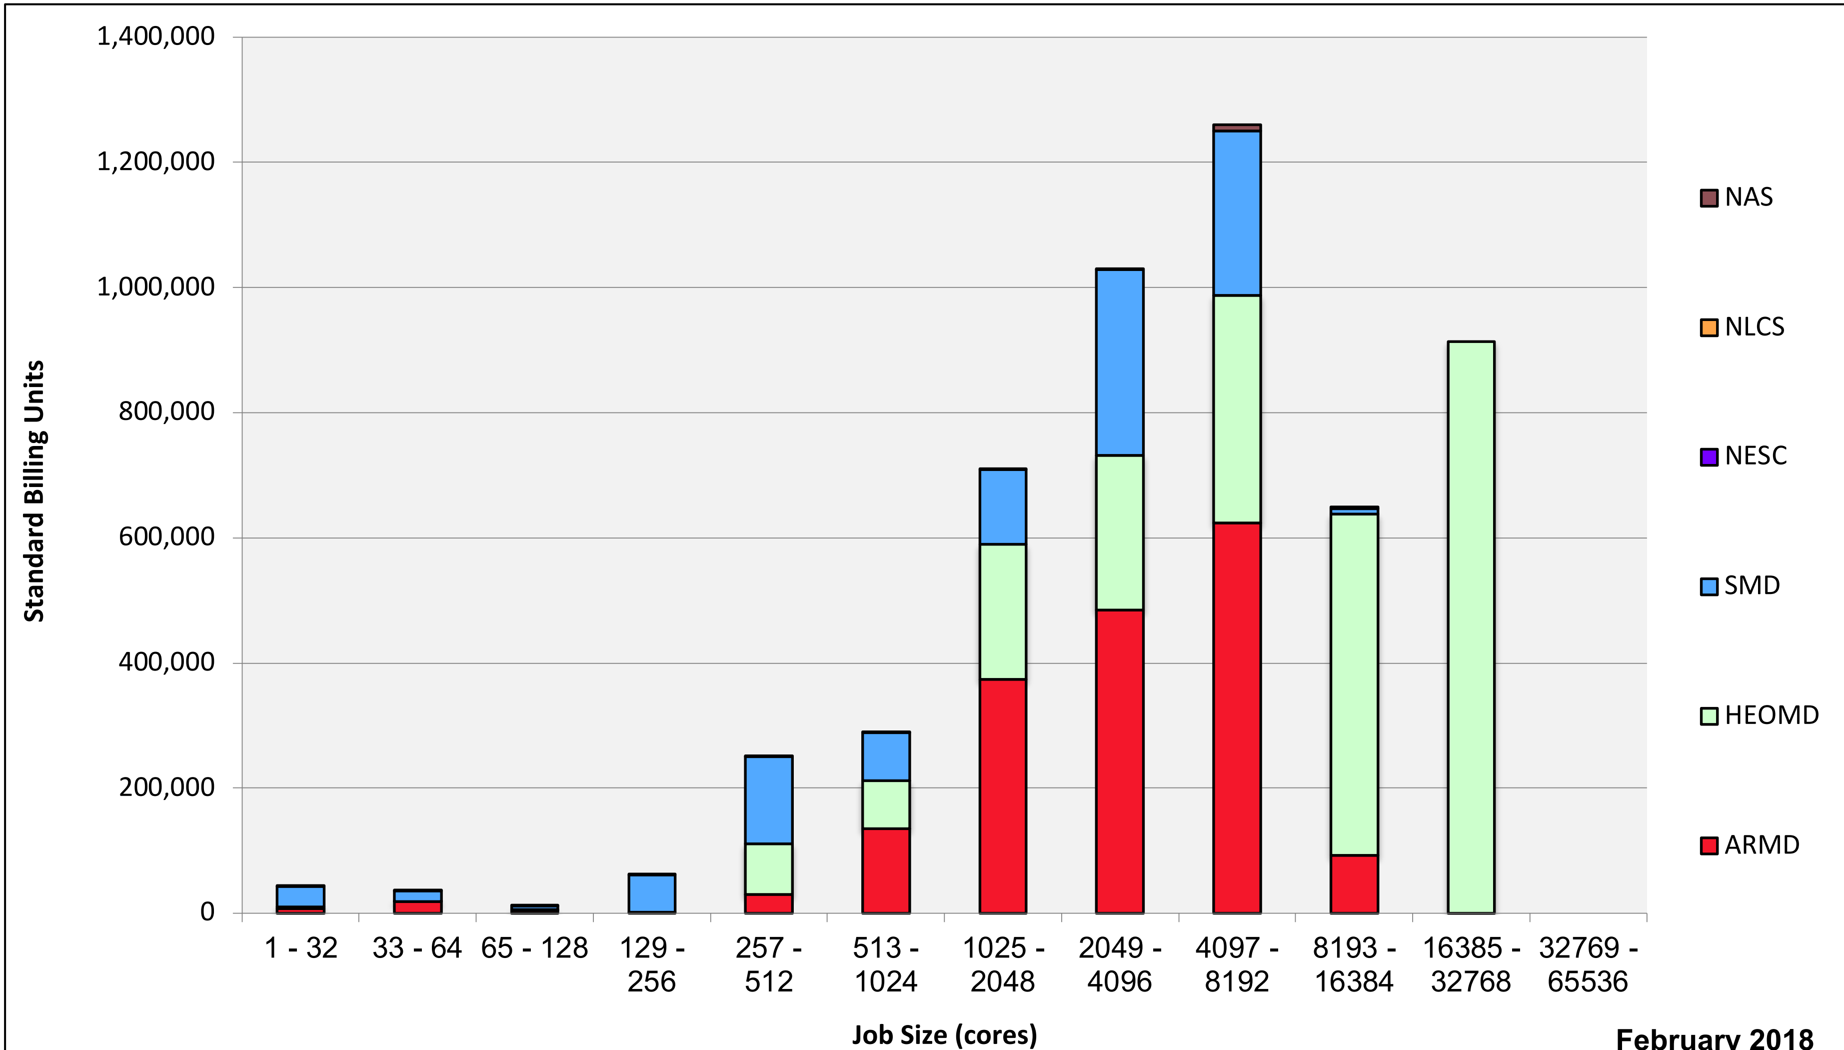

Electra:
Monthly Utilization by Size and Length

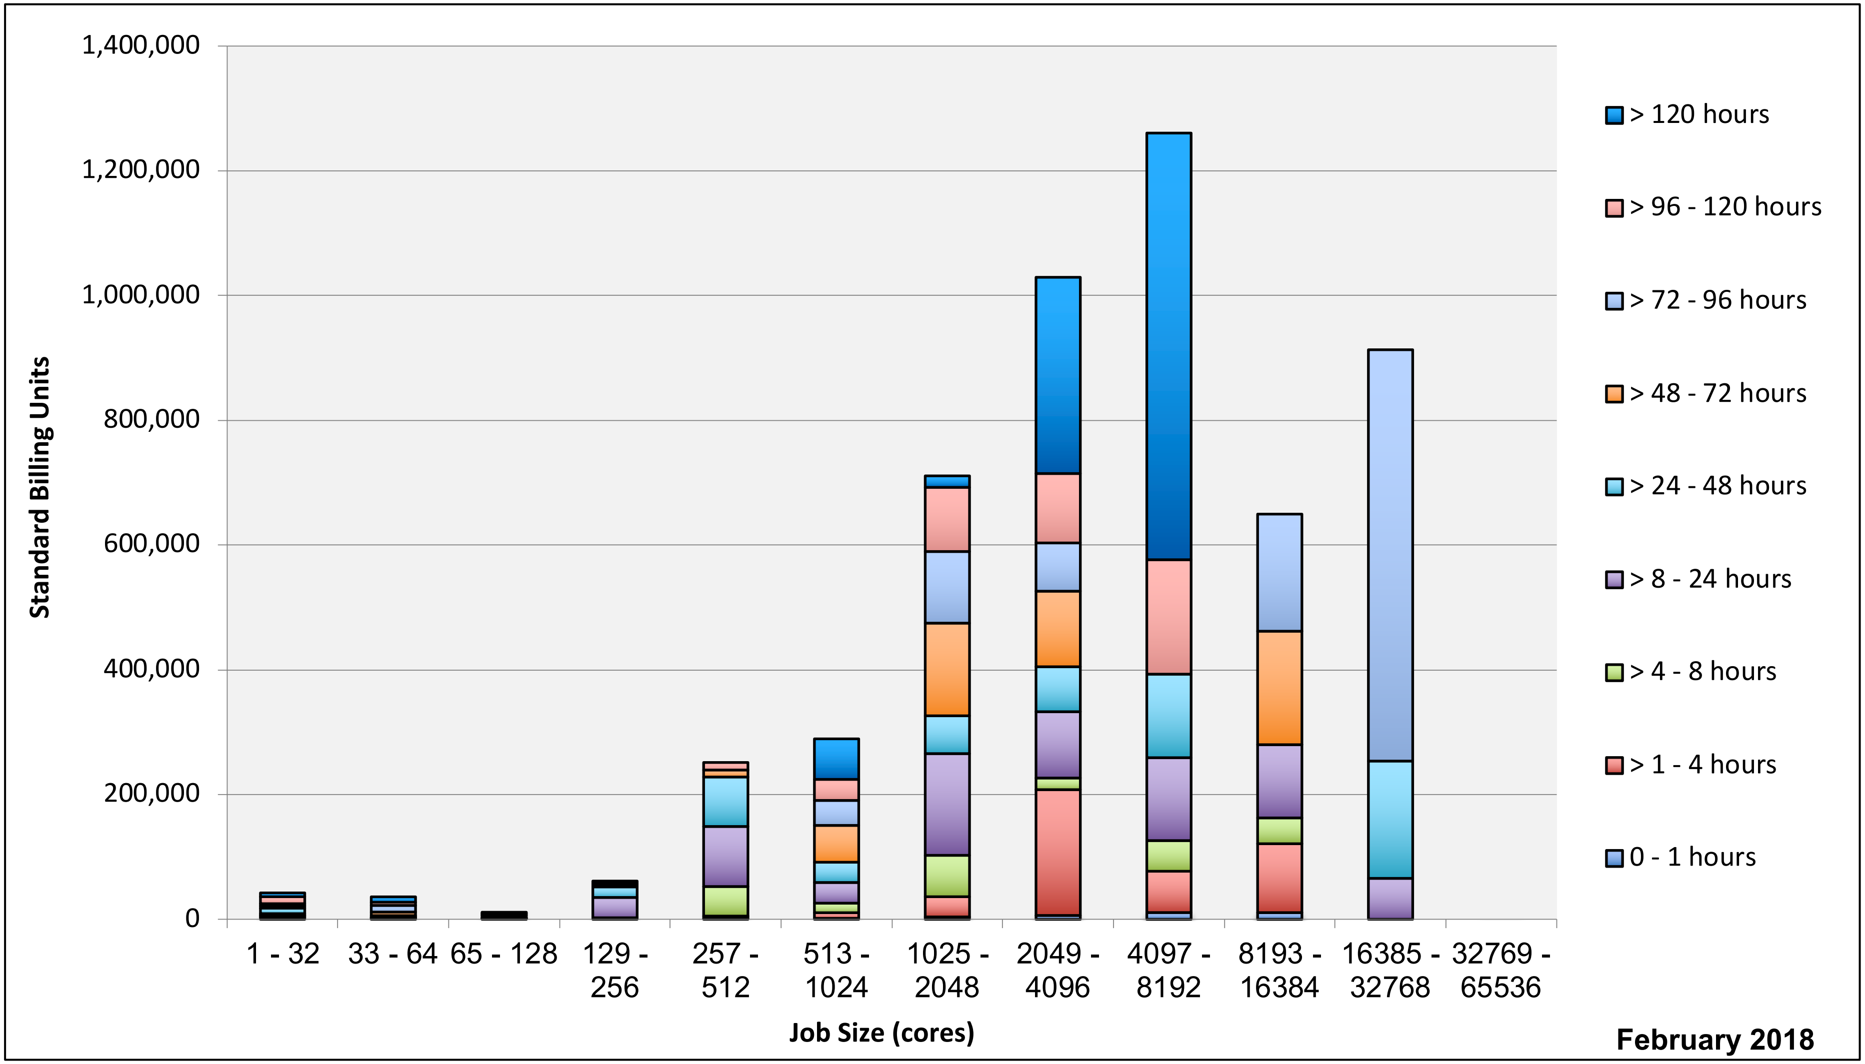

Electra:
Average Time to Clear All Jobs

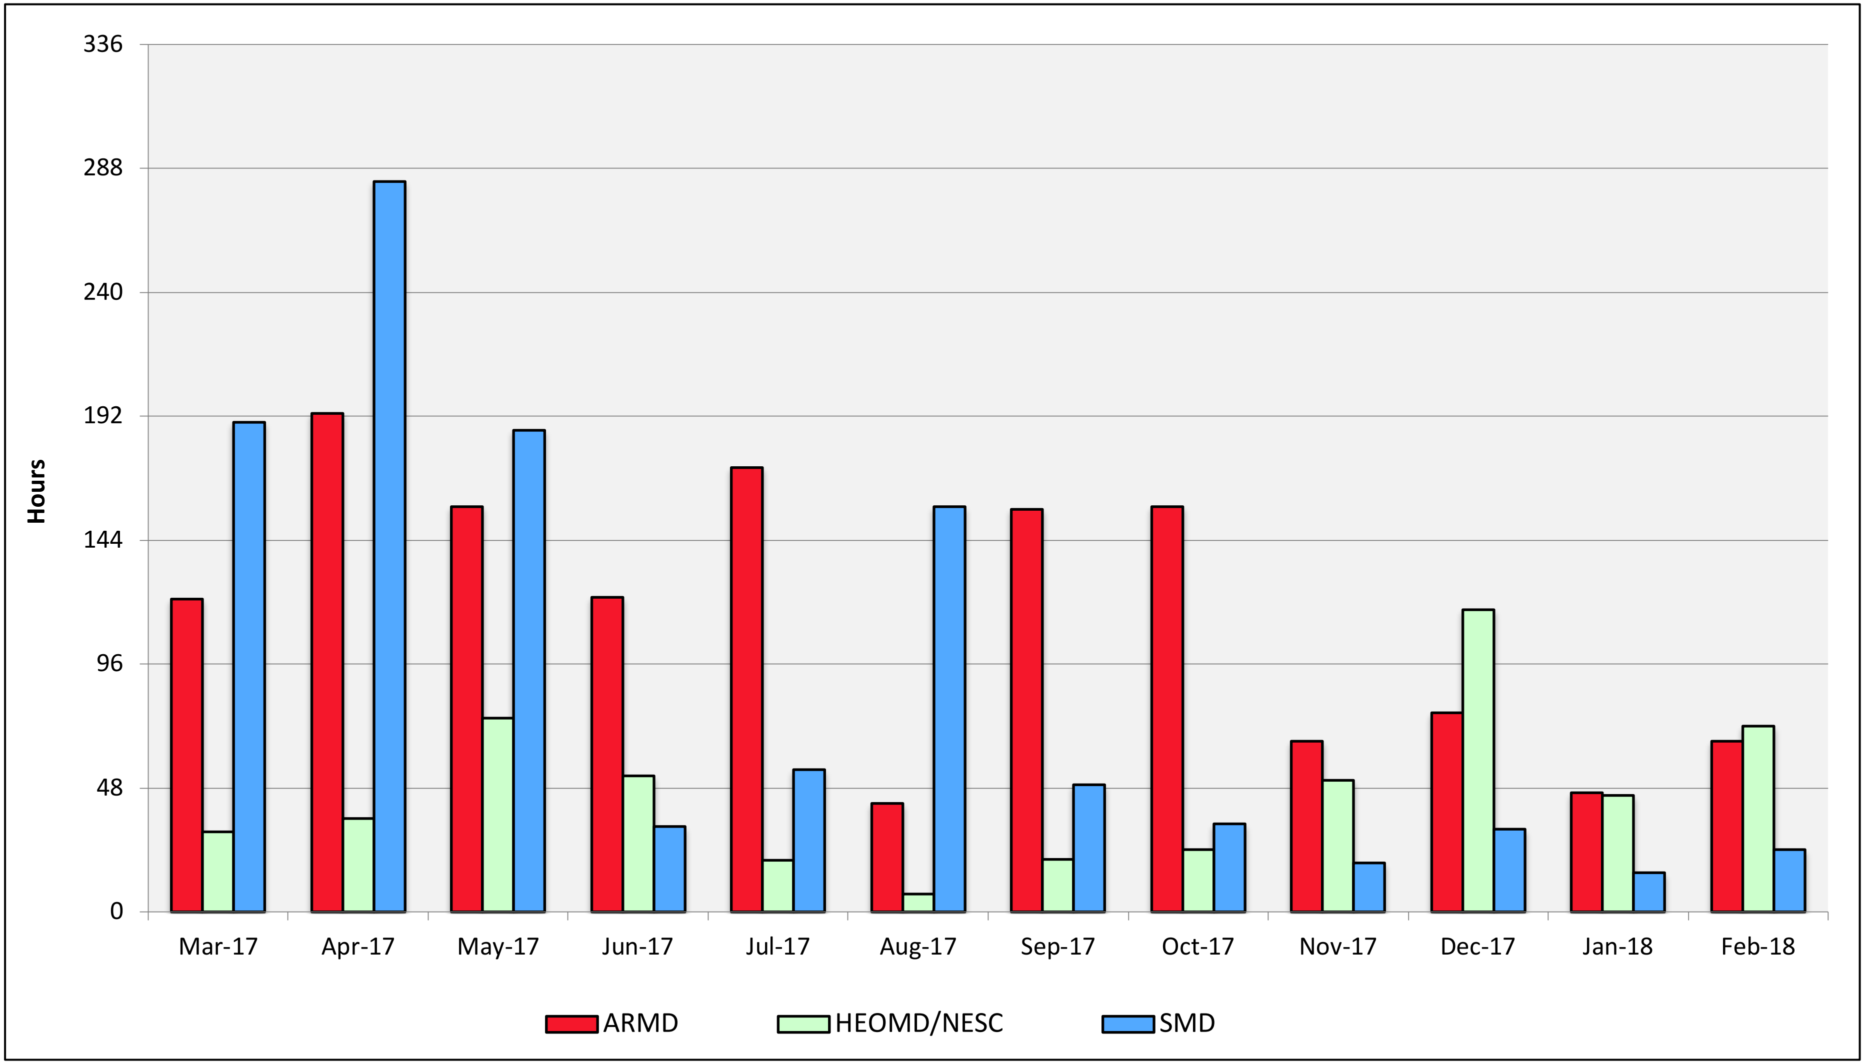

Electra:
Average Expansion Factor

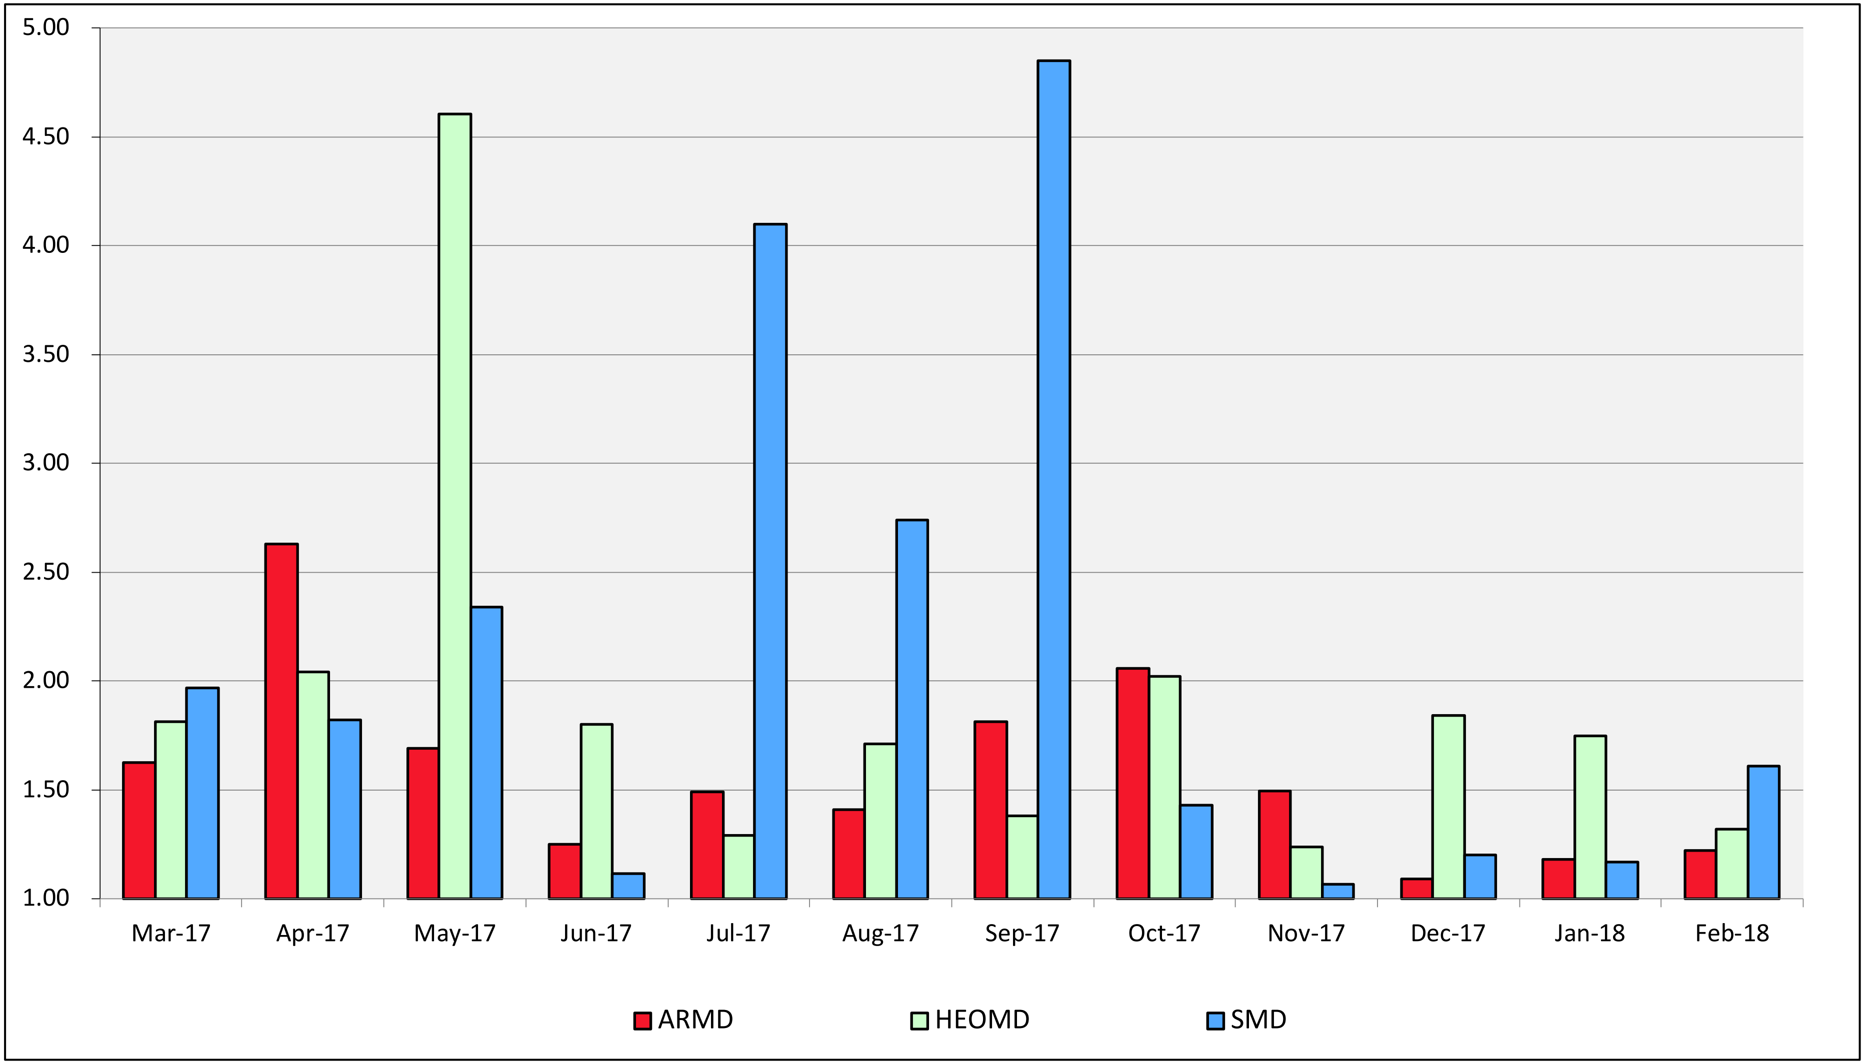

Merope: SBUs Reported, Normalized to 30-Day Month

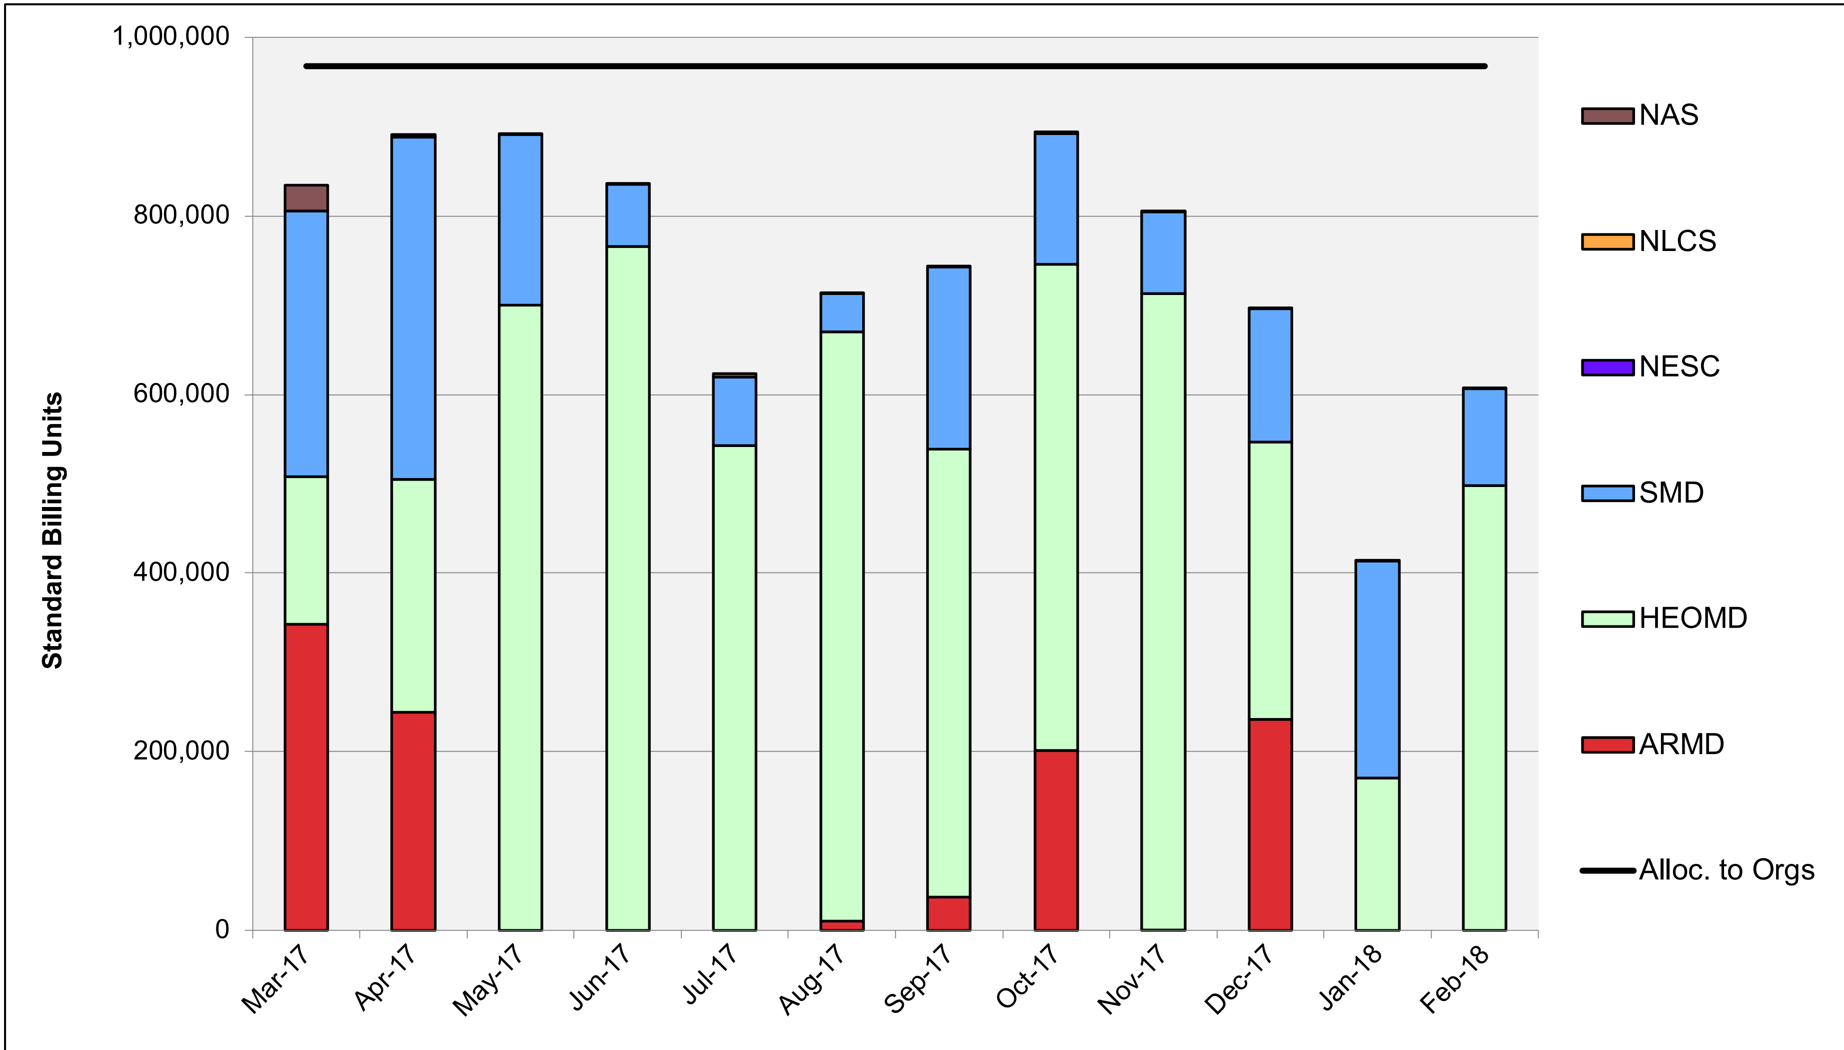

Merope:
Monthly Utilization by Job Length

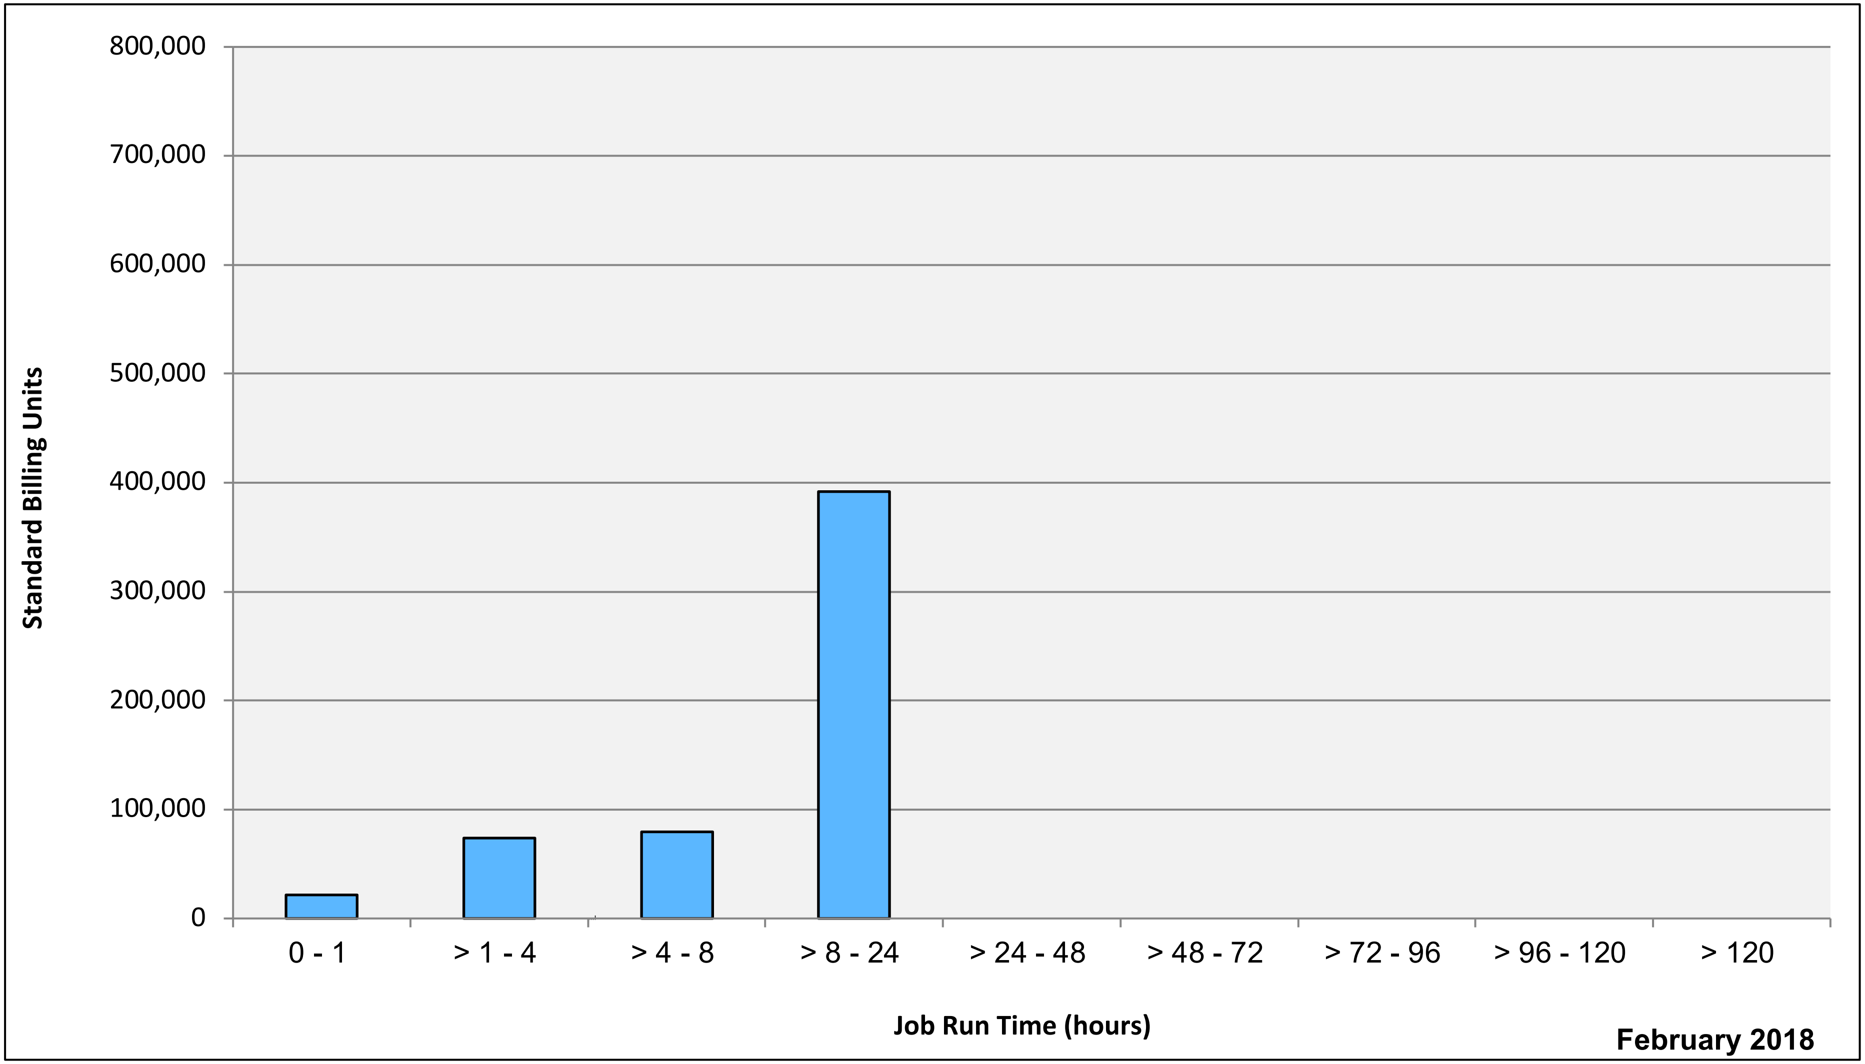

Merope: Monthly Utilization by Size and Mission

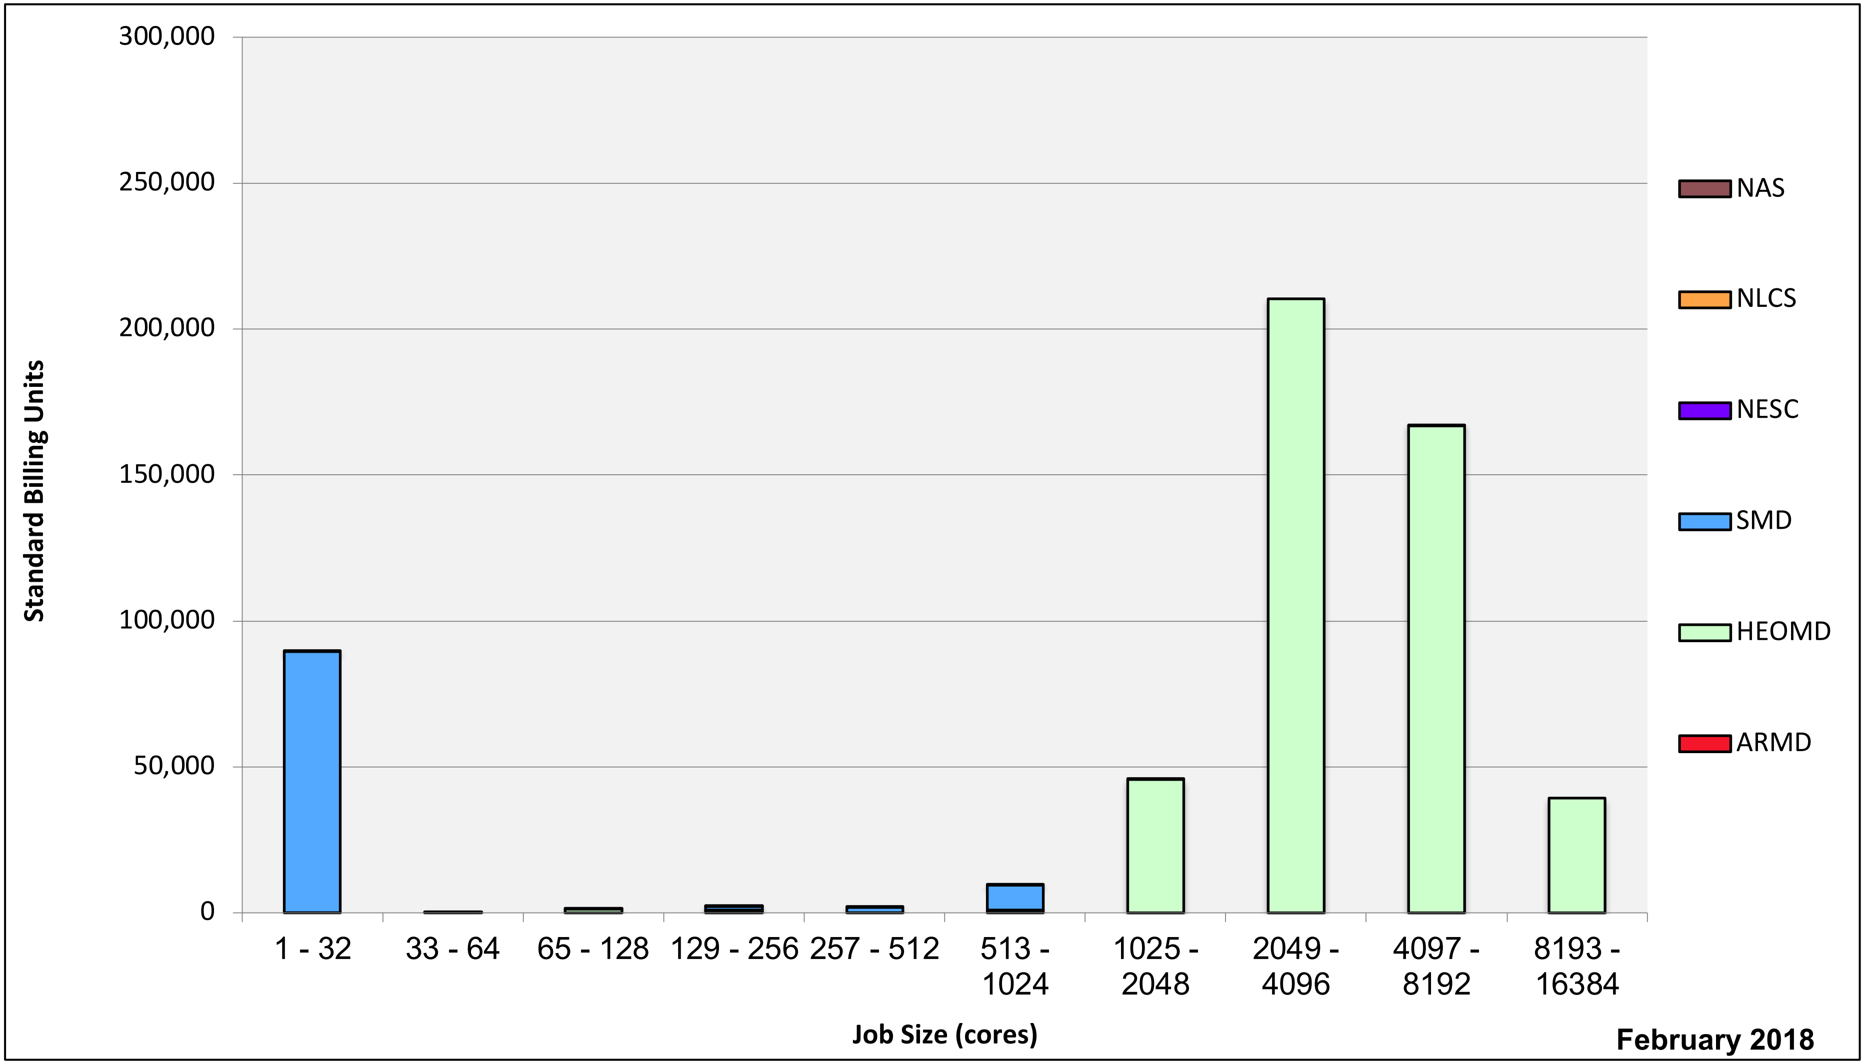

Merope: Monthly Utilization by Size and Length

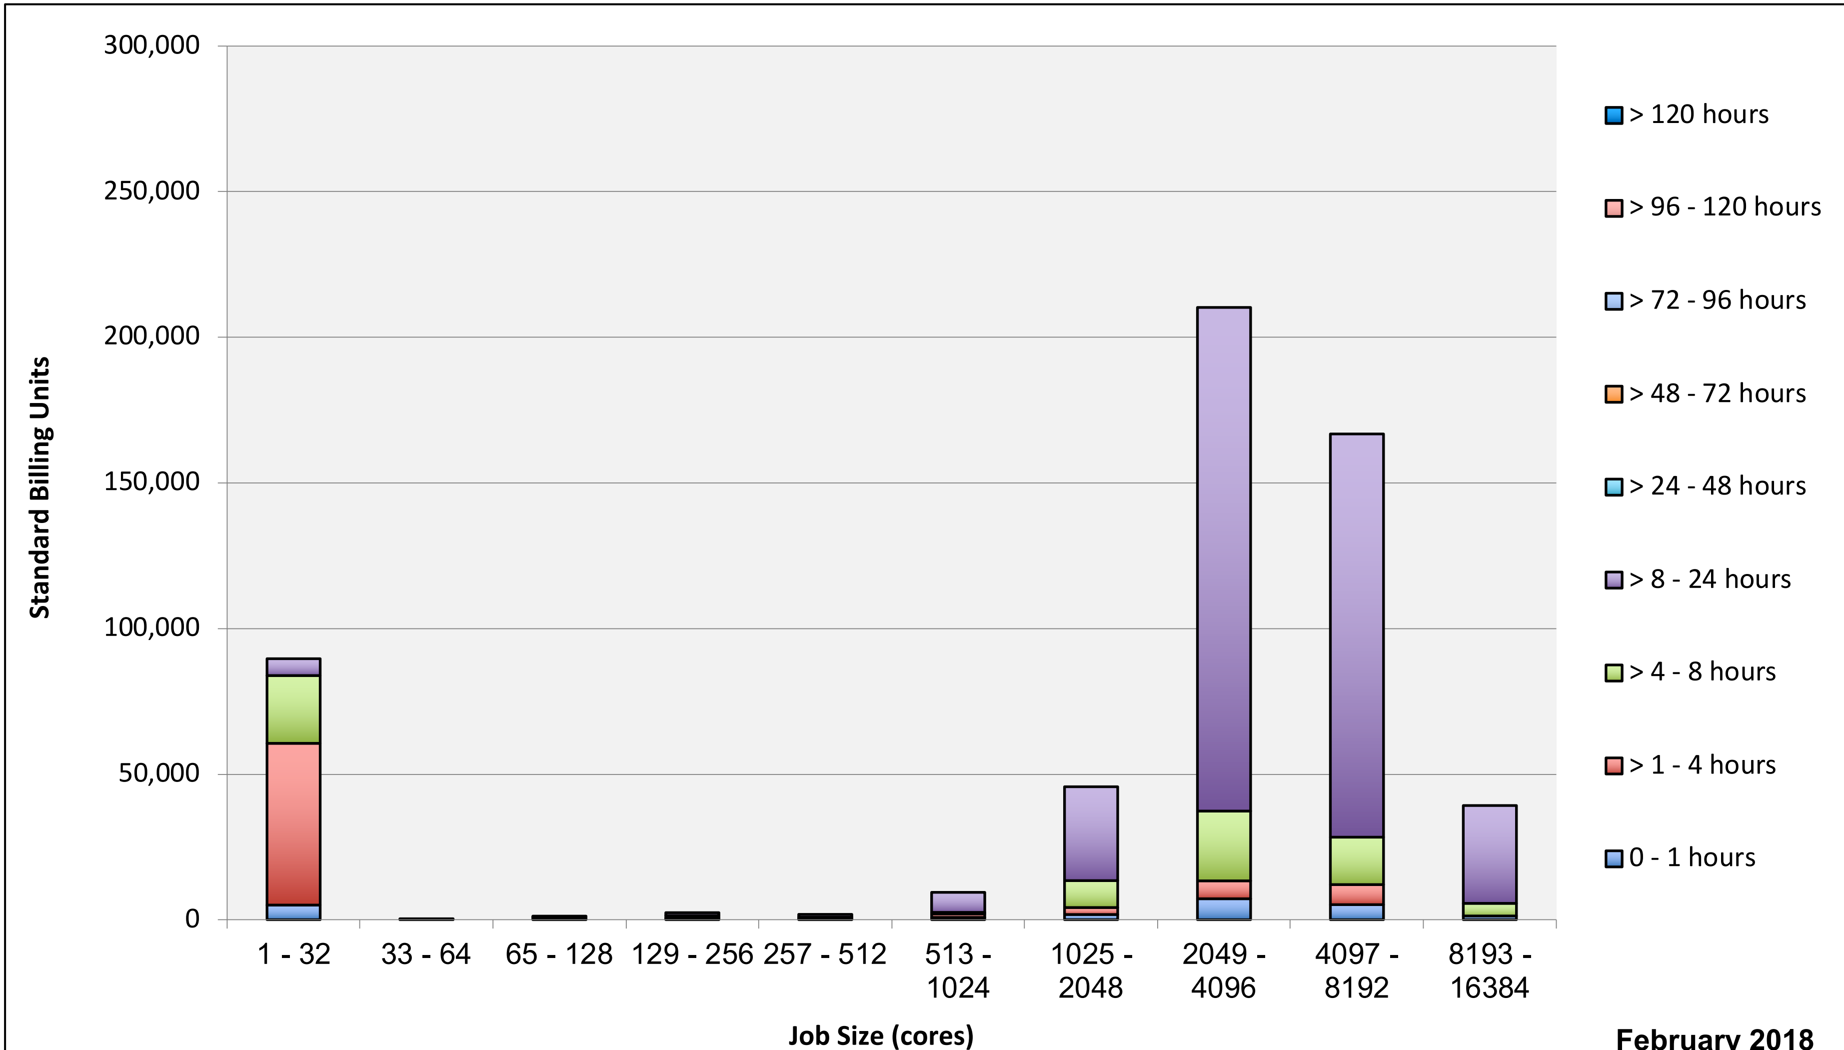

Merope:
Average Time to Clear All Jobs

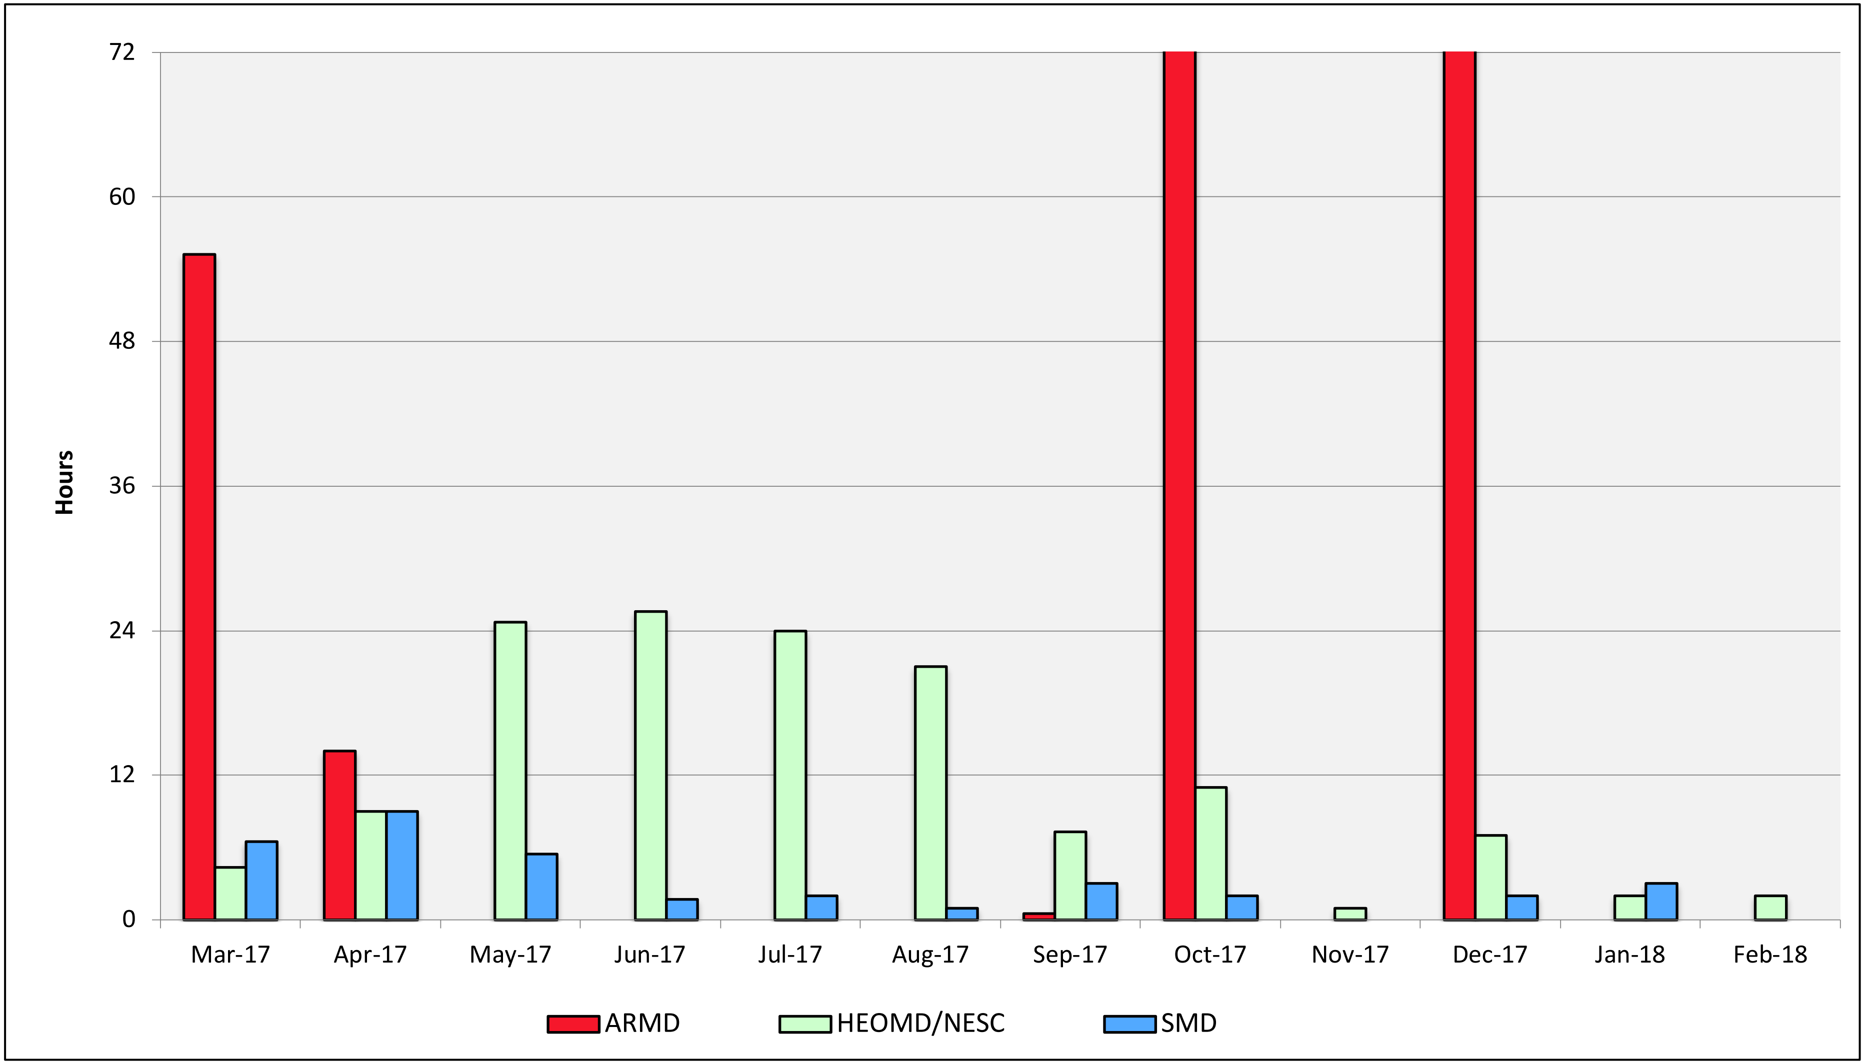

Merope:
Average Expansion Factor

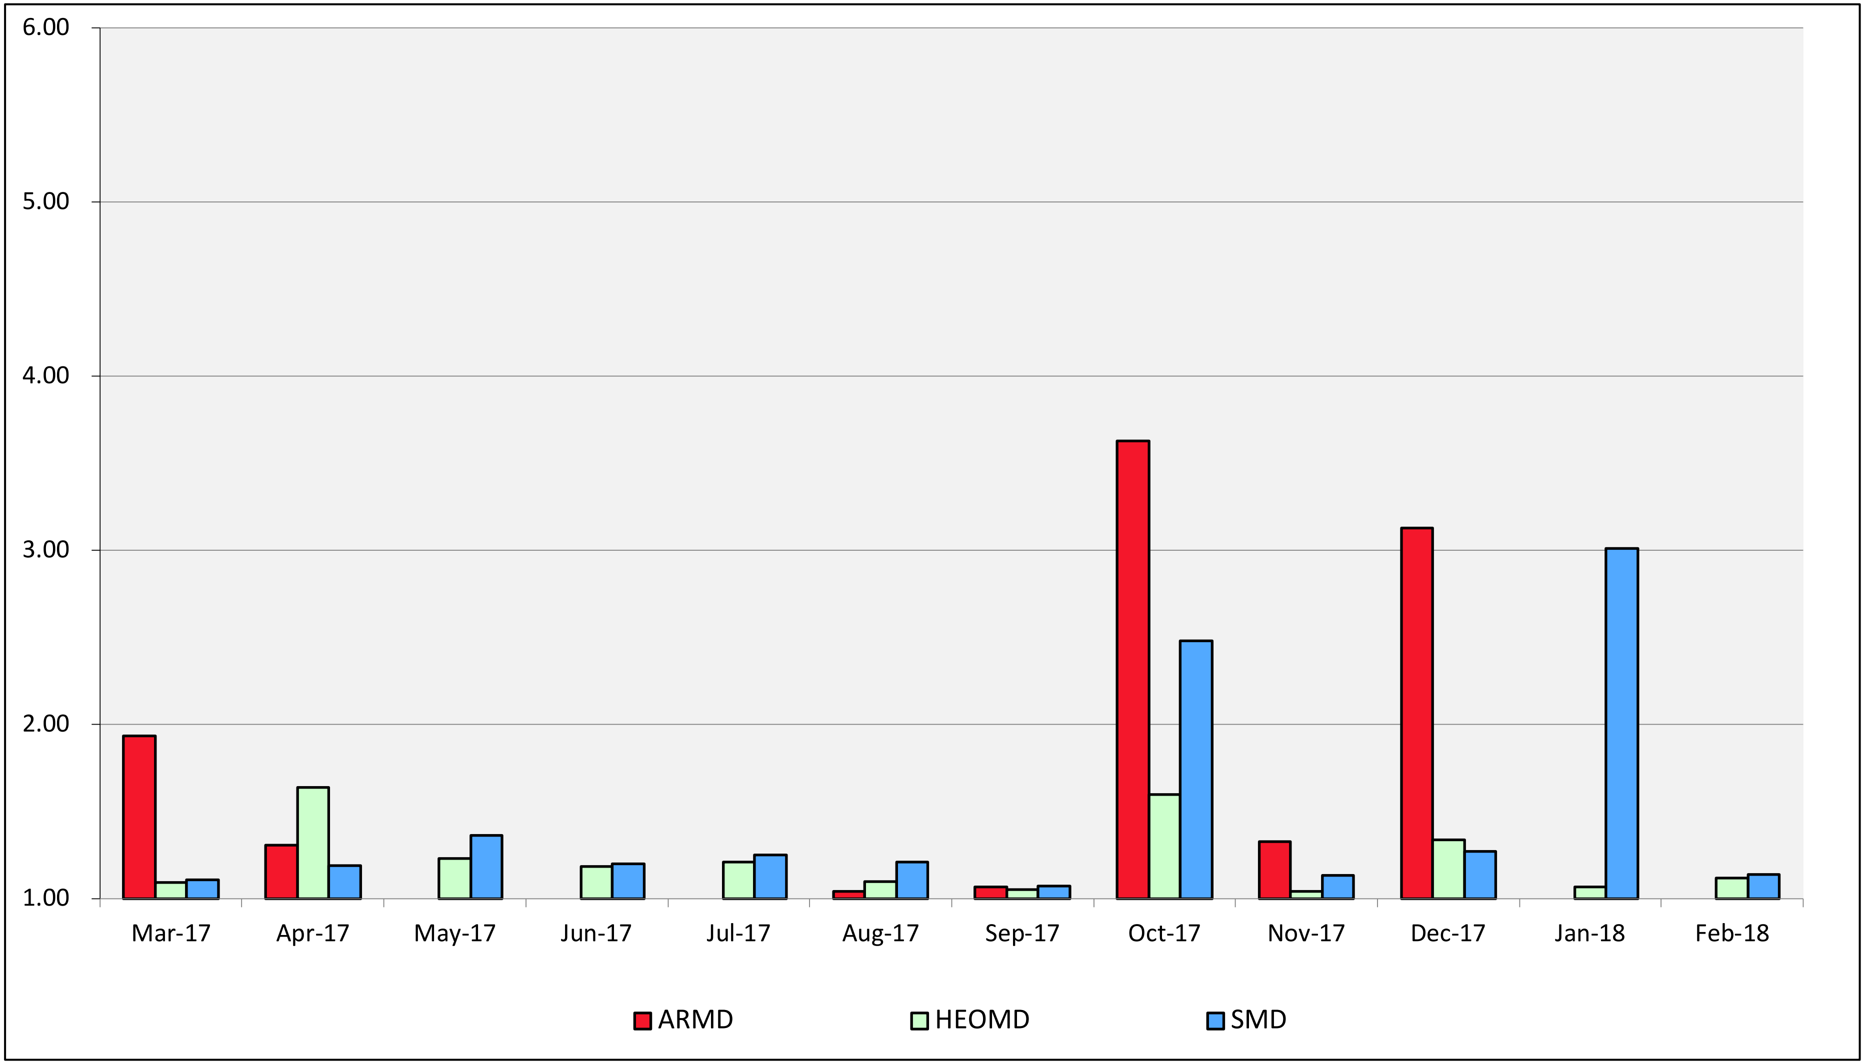

Endeavour:
SBUs Reported, Normalized to 30-Day Month

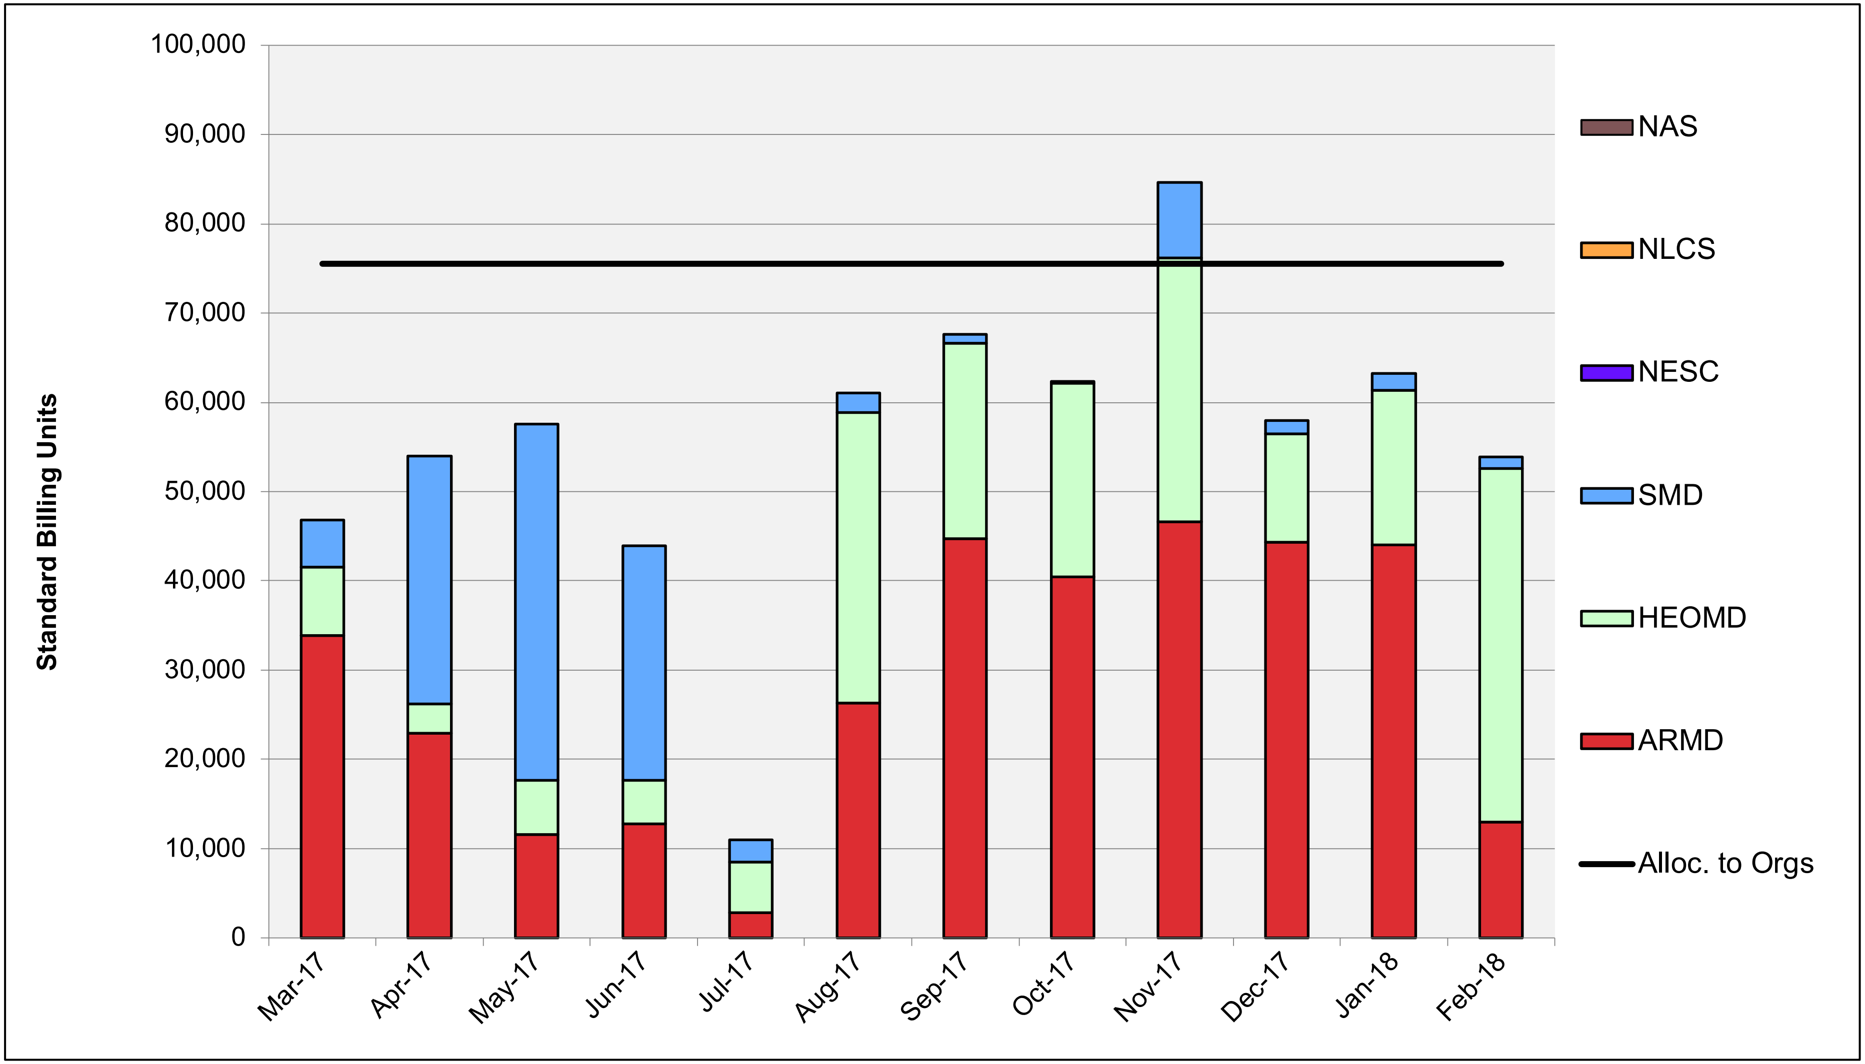

Endeavour:
Monthly Utilization by Job Length

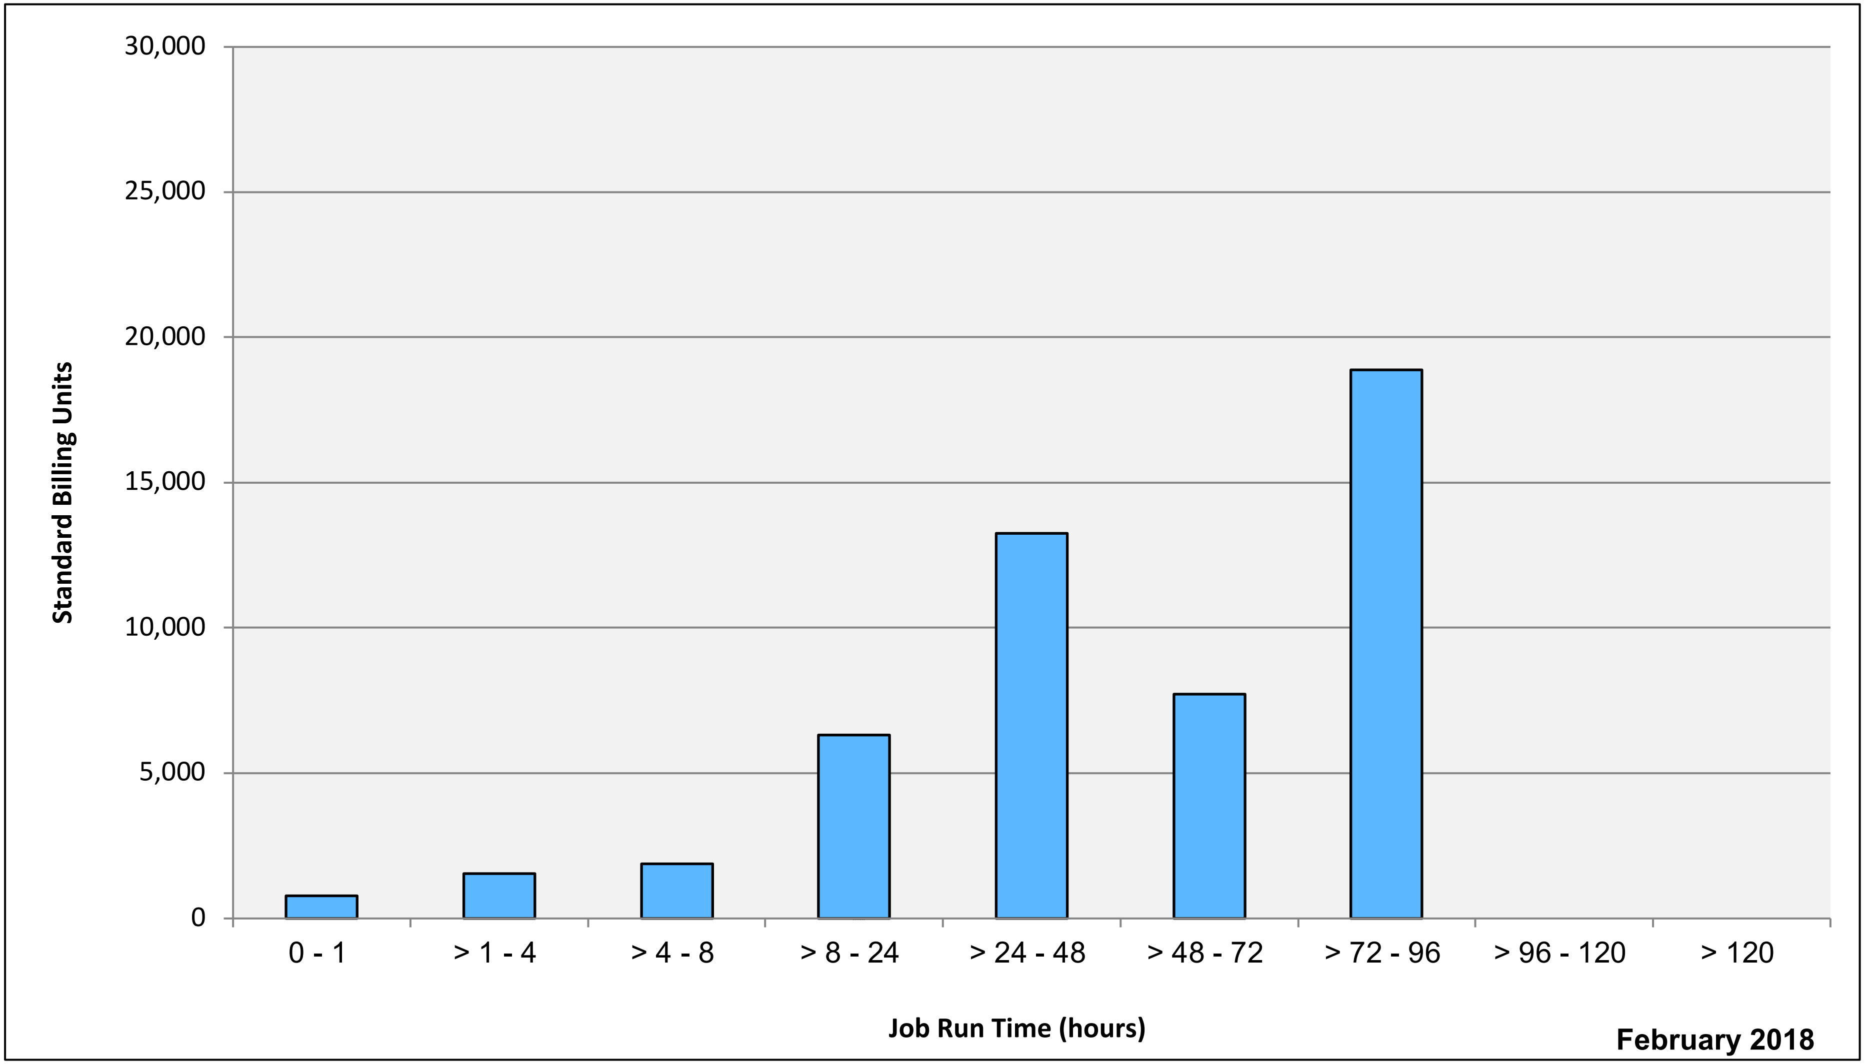

Endeavour:
Monthly Utilization by Size and Mission

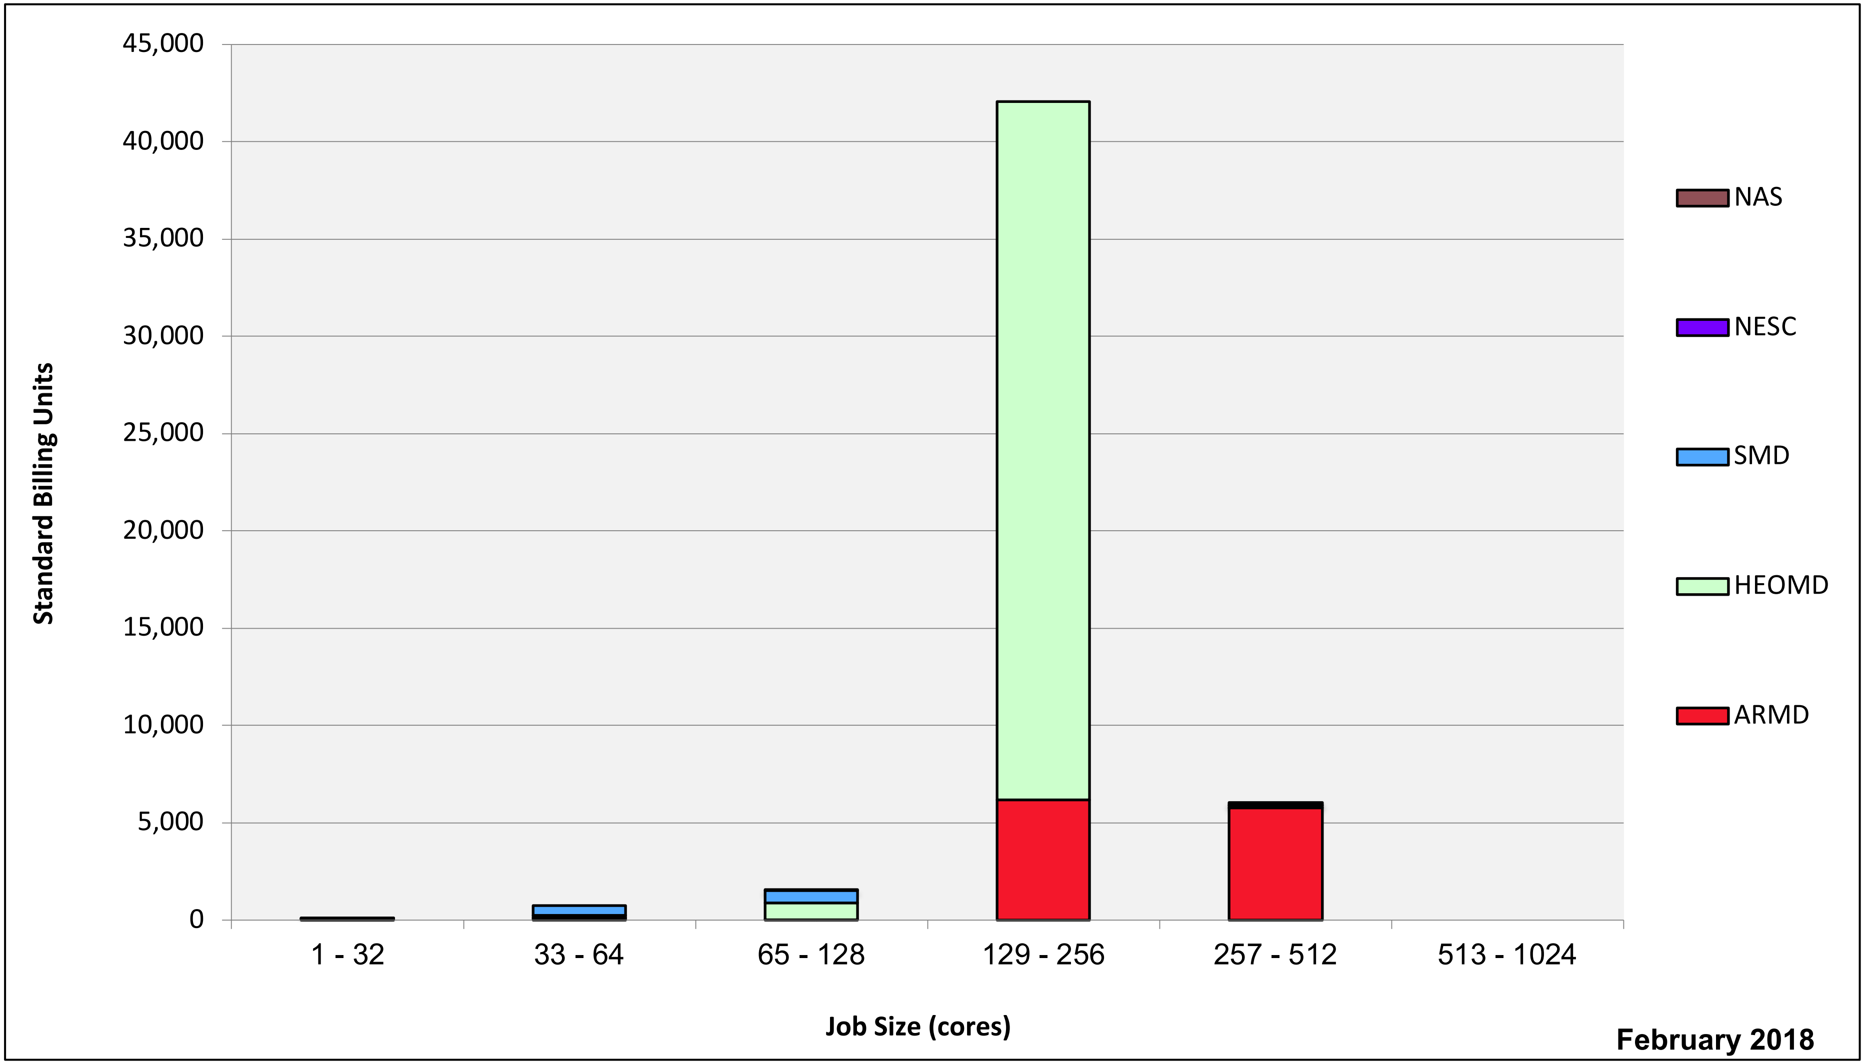

Endeavour:
Monthly Utilization by Size and Length

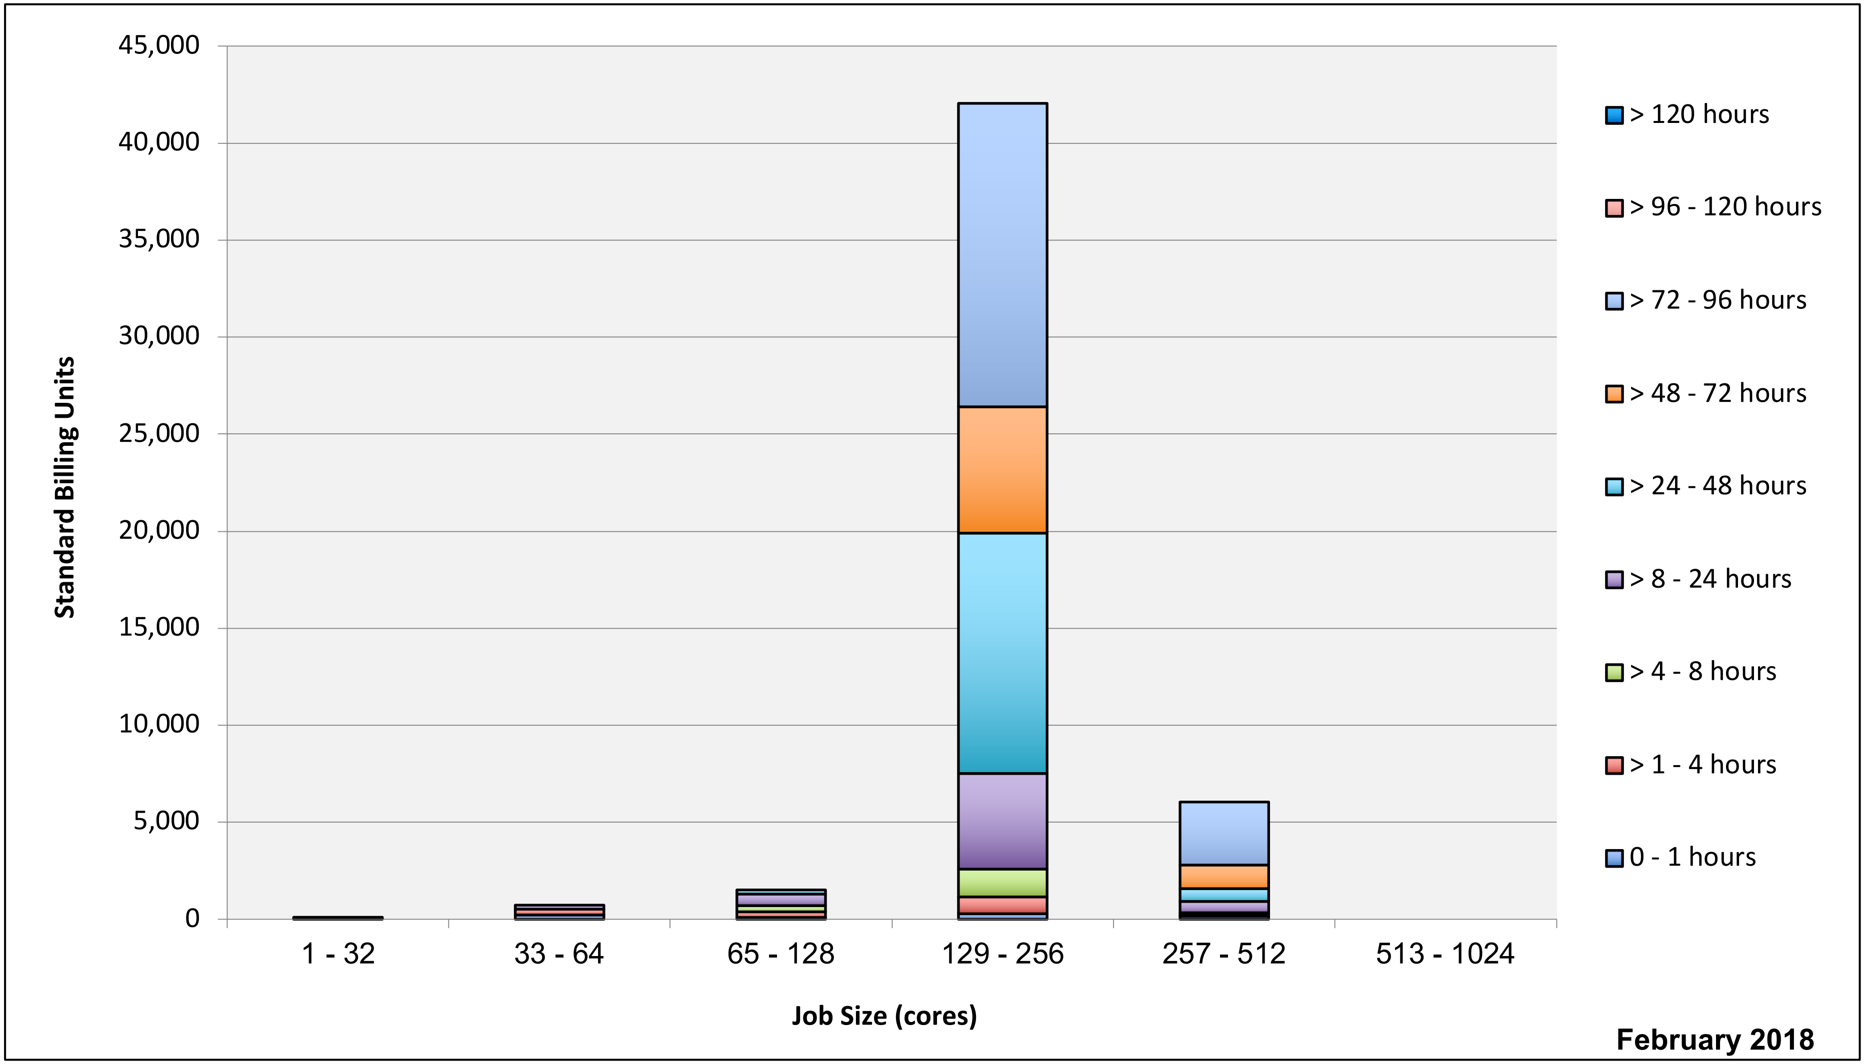

Endeavour: Average Time to Clear All Jobs

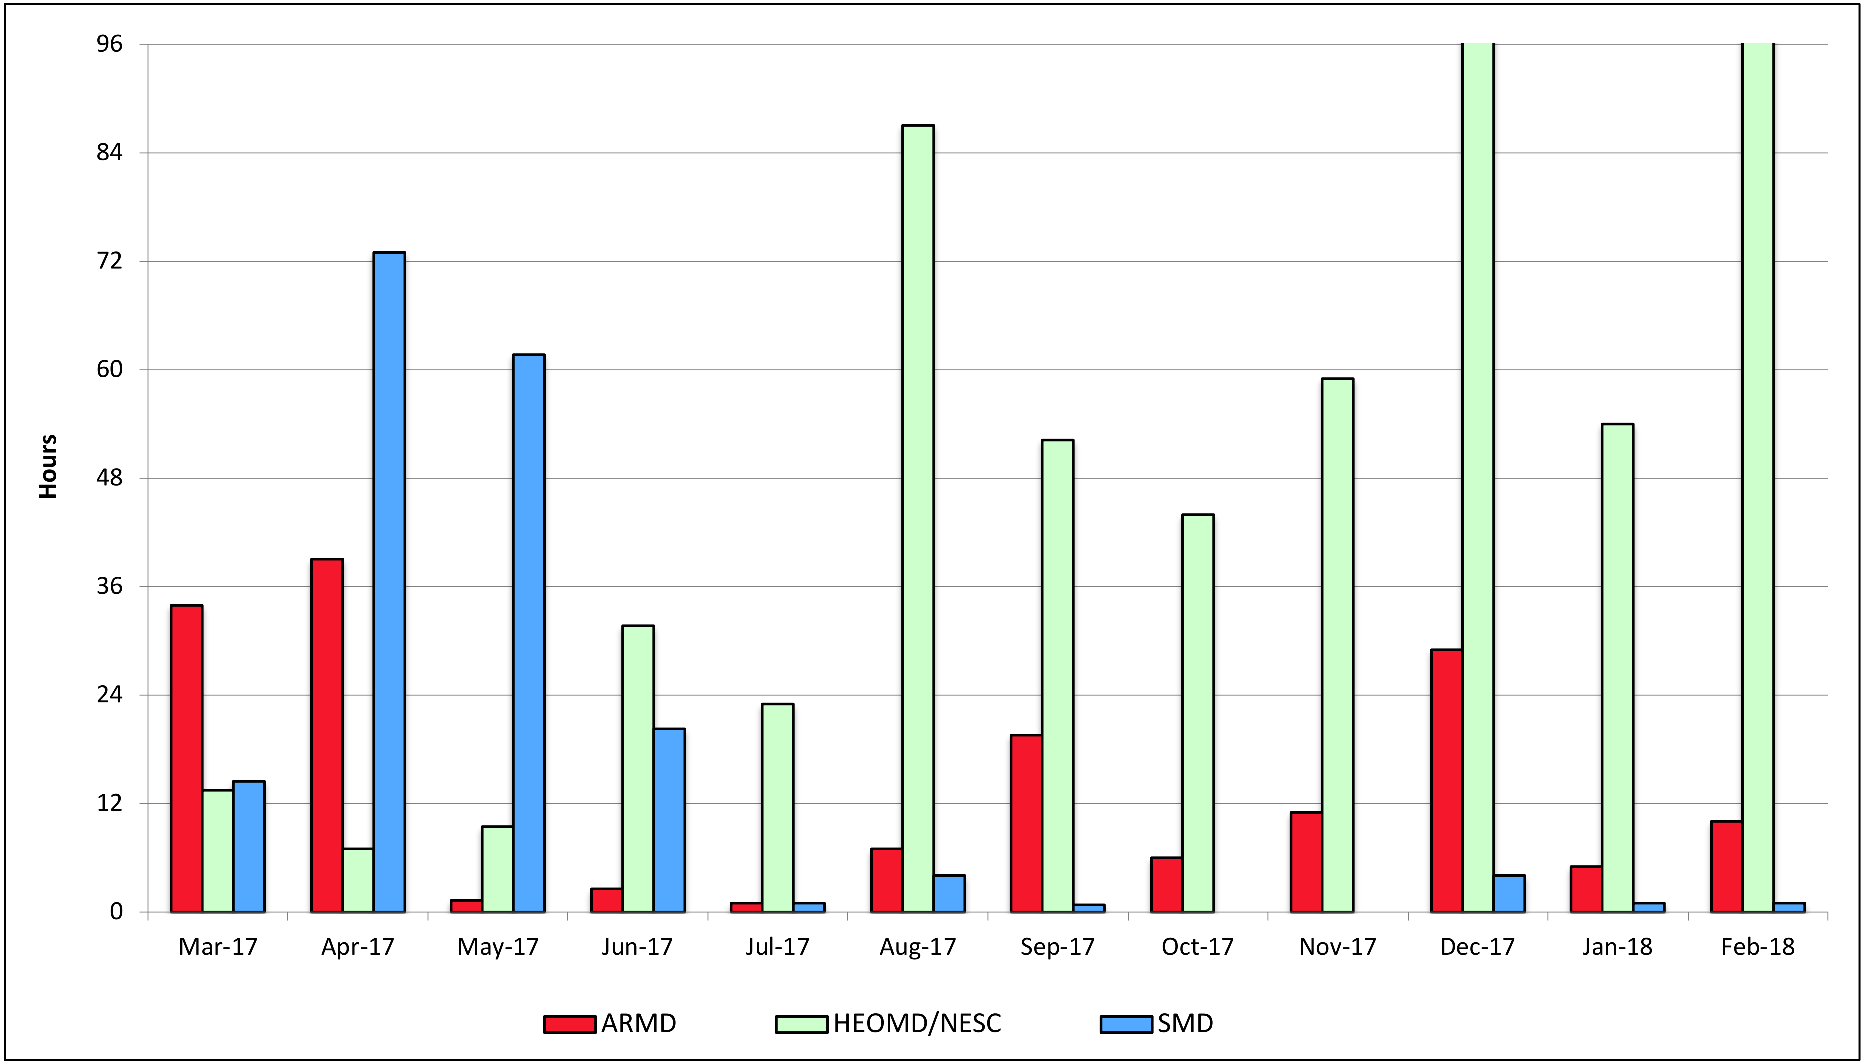

Endeavour: Average Expansion Factor